




Bayesian Conjugacy in Probit, Tobit, Multinomial Probit and Extensions: A Review and New Results

Niccolò Anceschi, Augusto Fasano, Daniele Durante & Giacomo Zanella


To cite this article: Niccolò Anceschi, Augusto Fasano, Daniele Durante & Giacomo Zanella (2023) Bayesian Conjugacy in Probit, Tobit, Multinomial Probit and Extensions: A Review and New Results, Journal of the American Statistical Association, 118:542, 1451-1469, DOI: 10.1080/01621459.2023.2169150

To link to this article: <https://doi.org/10.1080/01621459.2023.2169150>

 View supplementary material [↗](#)

 Published online: 03 Mar 2023.

 Submit your article to this journal [↗](#)

 Article views: 923

 View related articles [↗](#)

 View Crossmark data [↗](#)



Bayesian Conjugacy in Probit, Tobit, Multinomial Probit and Extensions: A Review and New Results

Niccolò Anceschi^a, Augusto Fasano^b, Daniele Durante^a, and Giacomo Zanella^a

^aDepartment of Decision Sciences and Bocconi Institute for Data Science and Analytics, Bocconi University, Milan, Italy; ^bCollegio Carlo Alberto, Turin, Italy

ABSTRACT

A broad class of models that routinely appear in several fields can be expressed as partially or fully discretized Gaussian linear regressions. Besides including classical Gaussian response settings, this class also encompasses probit, multinomial probit and tobit regression, among others, thereby yielding one of the most widely-implemented families of models in routine applications. The relevance of such representations has stimulated decades of research in the Bayesian field, mostly motivated by the fact that, unlike for Gaussian linear regression, the posterior distribution induced by such models does not seem to belong to a known class, under the commonly assumed Gaussian priors for the coefficients. This has motivated several solutions for posterior inference relying either on sampling-based strategies or on deterministic approximations that, however, still experience computational and accuracy issues, especially in high dimensions. The scope of this article is to review, unify and extend recent advances in Bayesian inference and computation for this core class of models. To address such a goal, we prove that the likelihoods induced by these formulations share a common analytical structure implying conjugacy with a broad class of distributions, namely the unified skew-normal (SUN), that generalize Gaussians to include skewness. This result unifies and extends recent conjugacy properties for specific models within the class analyzed, and opens new avenues for improved posterior inference, under a broader class of formulations and priors, via novel closed-form expressions, iid samplers from the exact SUN posteriors, and more accurate and scalable approximations from variational Bayes and expectation-propagation. Such advantages are illustrated in simulations and are expected to facilitate the routine-use of these core Bayesian models, while providing novel frameworks for studying theoretical properties and developing future extensions. Supplementary materials for this article are available online.

ARTICLE HISTORY

Received April 2022
Accepted December 2022

KEYWORDS

Bayesian computation; Data augmentation; Expectation-propagation; Truncated normal distribution; Unified skew-normal distribution; Variational Bayes

1. Introduction

The scope of this contribution is to review, unify, compare and extend both past and recent developments in Bayesian inference for probit (Bliss 1934), multinomial probit (Hausman and Wise 1978; Tutz 1991; Stern 1992), and tobit (Tobin 1958) regression models, along with related extensions to multivariate, skewed, nonlinear, and dynamic contexts. Although such models are core formulations in statistics (DeMaris 2004; Greene 2008; Agresti 2013) and often appear as building-blocks in more complex constructions (e.g., Chipman, George, and McCulloch 2010; Rodriguez and Dunson 2011), Bayesian inference under the associated likelihoods still presents open challenges that have motivated decades of active research in the field (Chopin and Ridgway 2017). This is mainly due to the presence in the likelihood of Gaussian cumulative distribution functions arising from a partial or full discretization of a set of continuous latent utilities under a discrete choice perspective (Greene 2008), thus hindering Gaussian conjugacy when combined with the commonly-assumed multivariate normal priors for the coefficients in β .

This lack of conjugacy for such a routinely implemented class of models motivates ongoing efforts to develop effective MCMC-based sampling methods and accurate deterministic approximations of the posterior distribution to perform Bayesian inference in probit (Albert and Chib 1993; Holmes and Held 2006; Consonni and Marin 2007; Chopin and Ridgway 2017), tobit (Chib 1992; Chib, Greenberg, and Jeliazkov 2009; Loaiza-Maya et al. 2022), multinomial probit (Albert and Chib 1993; McCulloch and Rossi 1994; Nobile 1998; McCulloch, Polson, and Rossi 2000; Albert and Chib 2001; Imai and Van Dyk 2005; Loaiza-Maya and Nibbering 2022), and their multivariate, skewed, dynamic and nonlinear generalizations (Chib and Greenberg 1998; Chen, Dey, and Shao 1999; Andrieu and Doucet 2002; Sahu, Dey, and Branco 2003; Kuss, Rasmussen, and Herbrich 2005; Girolami and Rogers 2006; Bazán, Bolfarine, and Branco 2010; Talhouk, Doucet, and Murphy 2012; Soyer and Sung 2013; Riihimäki, Jylänki, and Vehtari 2014). Although these methods yield state-of-the-art implementations, there are still open questions on computational scalability, approximation accuracy and mixing, especially in high dimensions (Chopin and Ridgway

2017). Such issues, combined with the recent conjugacy results for probit models in Durante (2019), have led to renewed interest in closed-form solutions for Bayesian inference under these formulations. In particular, Durante (2019) proved that the posterior for the β coefficients in Bayesian probit regression under Gaussian priors belongs to the class of unified skew-normal (SUN) distributions (Arellano-Valle and Azzalini 2006) and, more generally, that SUNs are conjugate to probit regression. The SUN class extends multivariate Gaussians to include skewness, and its analytical properties have led to rapid extensions of the original results in Durante (2019) to multinomial probit (Fasano and Durante 2022), dynamic multivariate probit (Fasano et al. 2021), Gaussian processes (Cao, Durante, and Genton 2022), skewed Gaussian processes (Benavoli, Azzimonti, and Piga 2020, 2021), skew-elliptical link functions (Zhang et al. *in press*) and rounded data (Kowal 2022), while facilitating the development of more accurate approximations (Fasano, Durante, and Zanella 2022).

These advancements provide yet unexplored opportunities for Bayesian inference under such models via novel closed-form expressions, tractable Monte Carlo methods based on iid samples from the exact SUN posteriors, and more accurate and scalable approximations from variational Bayes (VB) (e.g., Blei, Kucukelbir, and McAuliffe 2017) and expectation-propagation (EP) (e.g., Chopin and Ridgway 2017). However, most of these new developments focus on specific sub-classes of models within a potentially broader family of formulations that rely on partially or fully discretized Gaussian latent utilities. Therefore, there is still the lack of a unified framework which would be practically and conceptually useful to derive general conjugacy results along with broadly applicable closed-form solutions, Monte Carlo methods and improved approximations of the posterior distribution. For instance, conjugacy results for tobit models (Tobin 1958) are yet missing in the literature, however, as it will be clarified in Section 3, SUNs are conjugate also to this class. Such a comprehensive treatment would also help to clarify these advancements in the light of previously-developed state-of-the-art MCMC methods and approximations, and would serve as a catalyst of applied, methodological and theoretical research to further expand the set of solutions for this broad class of models.

This article aims at covering the aforementioned gap to boost the routine use of these core Bayesian models, and provide comprehensive frameworks for studying general theoretical properties and developing future extensions. As a first step toward accomplishing this goal, Section 2 unifies probit, tobit, multinomial probit and related extensions by reformulating the associated likelihoods as special cases of a general form that relies on the product between multivariate Gaussian densities and cumulative distributions, both evaluated at a linear combination of the coefficients β . Such a unified formulation is crucial to prove a general result in Section 3 which states that SUN distributions are conjugate priors for any model whose likelihood can be expressed as a special case of the one defined in Section 2. This result unifies available findings for probit (Durante 2019), multinomial probit (Fasano and Durante 2022) and dynamic multivariate probit (Fasano et al. 2021), among others, while extending SUN conjugacy properties to a broader class of Bayesian models for which similar results have not

appeared yet in the literature. Notable examples are tobit (Tobin 1958), and any extension of probit, tobit and multinomial probit which replaces the Gaussian latent utilities with skew-normal ones (Chen, Dey, and Shao 1999; Sahu, Dey, and Branco 2003; Bazán, Bolfarine, and Branco 2010), among others; see Section S1 in the supplementary materials.

As discussed in Section 4, this unified conjugacy result is also practically relevant since it allows to inherit all the recent methodological and computational developments for Bayesian inference under SUN posteriors in probit and multinomial probit to the entire class of models presented in Section 2. These advancements include novel closed-form expressions for relevant posterior moments, marginal likelihoods and predictive distributions, along with improved Monte Carlo methods and deterministic approximations from VB and EP. These solutions are presented in detail within Section 4 along with a careful review of the previous state-of-the-art solutions recasted under the proposed general framework. An excellent overview of these previous strategies can be found in Chopin and Ridgway (2017), but the focus is on univariate probit models. Due to this, the present article will mostly consider the more recent developments relying on SUN conjugacy and on their discussion in the light of previous solutions, when adapted to the broader class of models and priors, beyond classical Bayesian probit regression. Consistent with this scope, Section 6 concludes with a general discussion that points toward several future research directions motivated by the unified framework developed in this article. Empirical studies illustrating the potentials of this unification are provided in Section 5. The proofs of the theoretical results in this article and an in-depth discussion of the computational costs for the methods presented in Section 4 can be found in the supplementary materials. R codes with tutorials to reproduce the analyses in Section 5 are available at <https://github.com/niccoloaanceschi/TobitSUN>.

2. A Unified Likelihood Representation

As discussed in Section 1, probit regression (Bliss 1934), tobit (Tobin 1958), multinomial probit (Hausman and Wise 1978; Tutz 1991; Stern 1992) and their extensions are core formulations within statistics, and, when viewed as specific examples of a more general representation which also includes classical Gaussian linear regression, arguably yield to one of the most widely-implemented classes of models in routine applications (DeMaris 2004; Greene 2008; Agresti 2013). In fact, all these formulations share a common generative construction, in that the corresponding responses can be defined as partially or fully discretized versions of continuous ones from a set of underlying Gaussian linear regressions (Chib 1992; Albert and Chib 1993; Chib and Greenberg 1998). In particular, let $z_i \in \mathbb{R}$ denote a latent continuous response available for each unit $i = 1, \dots, n$, and consider the standard linear regression model $z_i = \mathbf{x}_i^\top \beta + \varepsilon_i$, with iid noise $\varepsilon_i \sim N(0, \sigma^2)$, covariates' vector $\mathbf{x}_i = (x_{i1}, \dots, x_{ip})^\top$ and coefficients $\beta = (\beta_1, \dots, \beta_p)^\top$. Starting from this building-block formulation, general Gaussian linear regression models, probit models (Bliss 1934) and tobit regression (Tobin 1958) can be obtained by letting $y_i = z_i$, $y_i = \mathbb{1}(z_i > 0)$ and $y_i = \max\{z_i, 0\} = z_i \mathbb{1}(z_i > 0)$, respectively. The first

two constructions correspond to the limiting cases in which z_i is either entirely observed or dichotomized, respectively, whereas the third one represents the intermediate situation in which z_i is fully observed only if it exceeds value 0 (Chib 1992; Albert and Chib 1993). Multinomial probit regression (Hausman and Wise 1978; Tutz 1991; Stern 1992) for categorical responses $y_i \in \{1, \dots, L\}$ can be derived with a similar reasoning. For instance, in the formulation proposed by Stern (1992), the observed categorical response y_i is defined as $y_i = \operatorname{argmax}_l \{z_{i1}, \dots, z_{iL}\}$ where z_{i1}, \dots, z_{iL} are class-specific Gaussian latent utilities related to the covariates via a system of linear regressions $z_{il} = \mathbf{x}_i^\top \boldsymbol{\beta}_l + \varepsilon_{il}$ for each $l = 1, \dots, L$, with $\boldsymbol{\varepsilon}_i = (\varepsilon_{i1}, \dots, \varepsilon_{iL})^\top \sim N_L(\mathbf{0}, \boldsymbol{\Sigma})$.

As shown in Sections 2.1–2.4, these similarities in the generative models imply that the likelihoods induced by the above formulations and their extensions are all specific examples of the general form

$$p(\mathbf{y}|\boldsymbol{\beta}) = p(\bar{\mathbf{y}}_1|\boldsymbol{\beta})p(\bar{\mathbf{y}}_0|\boldsymbol{\beta}) \propto \phi_{\bar{n}_1}(\bar{\mathbf{y}}_1 - \bar{\mathbf{X}}_1\boldsymbol{\beta}; \bar{\boldsymbol{\Sigma}}_1)\Phi_{\bar{n}_0}(\bar{\mathbf{y}}_0 + \bar{\mathbf{X}}_0\boldsymbol{\beta}; \bar{\boldsymbol{\Sigma}}_0), \tag{1}$$

with $\phi_{\bar{n}_1}(\bar{\mathbf{y}}_1 - \bar{\mathbf{X}}_1\boldsymbol{\beta}; \bar{\boldsymbol{\Sigma}}_1)$ and $\Phi_{\bar{n}_0}(\bar{\mathbf{y}}_0 + \bar{\mathbf{X}}_0\boldsymbol{\beta}; \bar{\boldsymbol{\Sigma}}_0)$ denoting the density and the cumulative distribution function of the multivariate Gaussians $N_{\bar{n}_1}(\mathbf{0}, \bar{\boldsymbol{\Sigma}}_1)$ and $N_{\bar{n}_0}(\mathbf{0}, \bar{\boldsymbol{\Sigma}}_0)$, evaluated at $\bar{\mathbf{y}}_1 - \bar{\mathbf{X}}_1\boldsymbol{\beta}$ and $\bar{\mathbf{y}}_0 + \bar{\mathbf{X}}_0\boldsymbol{\beta}$, respectively, where $\bar{\mathbf{y}}_1 := \bar{\mathbf{y}}_1(\mathbf{y})$ and $\bar{\mathbf{y}}_0 := \bar{\mathbf{y}}_0(\mathbf{y})$ denote known response vectors obtained as a function of \mathbf{y} , whereas $\bar{\mathbf{X}}_1 := \bar{\mathbf{X}}_1(\mathbf{y}, \mathbf{X})$ and $\bar{\mathbf{X}}_0 := \bar{\mathbf{X}}_0(\mathbf{y}, \mathbf{X})$ are suitable matrices which can be directly derived from the original design matrix \mathbf{X} — whose rows comprise the vectors \mathbf{x}_i^\top , $i = 1, \dots, n$ — and, possibly, the responses in \mathbf{y} . Intuitively, (1) states that the likelihood for \mathbf{y} arises from a set of linear regressions for $\bar{n}_1 + \bar{n}_0$ Gaussian responses of which \bar{n}_1 are fully observed, thus contributing to the first factor in (1), whereas the remaining \bar{n}_0 are dichotomized and hence are captured by the second. For instance, the likelihood $\prod_{i=1}^n \Phi(\mathbf{x}_i^\top \boldsymbol{\beta})^{\mathbb{1}(y_i=1)} [1 - \Phi(\mathbf{x}_i^\top \boldsymbol{\beta})]^{\mathbb{1}(y_i=0)}$ under probit regression can be rewritten as $\prod_{i=1}^n \Phi[(2y_i - 1)\mathbf{x}_i^\top \boldsymbol{\beta}] = \Phi_n(\operatorname{diag}(2\mathbf{y} - \mathbf{1}_n)\mathbf{X}\boldsymbol{\beta}; \mathbf{I}_n)$, that coincides with (1) after letting $\bar{n}_1 = 0$, $\bar{n}_0 = n$, $\bar{\mathbf{y}}_0 = \mathbf{0}$, $\bar{\mathbf{X}}_0 = \operatorname{diag}(2\mathbf{y} - \mathbf{1}_n)\mathbf{X}$ and $\bar{\boldsymbol{\Sigma}}_0 = \mathbf{I}_n$. This example is a degenerate case where $\bar{n}_1 = 0$, meaning that there are no fully observed Gaussian responses and $\bar{n}_0 = n$ dichotomized ones; refer to Sections 2.1–2.4 for additional examples which further clarify (1). To simplify notation, in the following we write $\phi(\cdot)$ and $\Phi(\cdot)$ for denoting, respectively, the density and cumulative distribution function of a univariate standard Gaussian variable, with mean 0 and variance 1.

2.1. Linear Regression and Multivariate Linear Regression

Although the focus of this article is on models beyond the classical Gaussian response setting, it is worth emphasizing that also this class induces likelihoods which are special cases of (1). For instance, Gaussian linear regression $(y_i|\boldsymbol{\beta}) \sim N(\mathbf{x}_i^\top \boldsymbol{\beta}, \sigma^2)$, independently for $i = 1, \dots, n$, is directly recovered after noticing that the induced likelihood

$$p(\mathbf{y}|\boldsymbol{\beta}) = \prod_{i=1}^n \phi(y_i - \mathbf{x}_i^\top \boldsymbol{\beta}; \sigma^2) = \phi_n(\mathbf{y} - \mathbf{X}\boldsymbol{\beta}; \sigma^2\mathbf{I}_n), \tag{2}$$

coincides with (1), when letting $\bar{n}_0 = 0$, $\bar{n}_1 = n$, $\bar{\mathbf{y}}_1 = \mathbf{y}$, $\bar{\mathbf{X}}_1 = \mathbf{X}$ and $\bar{\boldsymbol{\Sigma}}_1 = \sigma^2\mathbf{I}_n$. As a consequence, also heteroscedastic

and correlated versions can be incorporated by replacing $\sigma^2\mathbf{I}_n$ with a general covariance matrix. Similarly, the likelihood associated with multivariate Gaussian response data from $(\mathbf{y}_i|\boldsymbol{\beta}) \sim N_m(\mathbf{X}_i\boldsymbol{\beta}, \boldsymbol{\Sigma})$, independently for $i = 1, \dots, n$, can be written as

$$p(\mathbf{y}|\boldsymbol{\beta}) = \prod_{i=1}^n \phi_m(\mathbf{y}_i - \mathbf{X}_i\boldsymbol{\beta}; \boldsymbol{\Sigma}) = \phi_{n \cdot m}(\mathbf{y} - \mathbf{X}\boldsymbol{\beta}; \mathbf{I}_n \otimes \boldsymbol{\Sigma}), \tag{3}$$

where $\mathbf{y} = (\mathbf{y}_1^\top, \dots, \mathbf{y}_n^\top)^\top$, $\mathbf{X} = (\mathbf{X}_1^\top, \dots, \mathbf{X}_n^\top)^\top$ and \otimes denotes the Kronecker product. Setting $\bar{n}_0 = 0$, $\bar{n}_1 = n \cdot m$, $\bar{\mathbf{y}}_1 = \mathbf{y}$, $\bar{\mathbf{X}}_1 = \mathbf{X}$ and $\bar{\boldsymbol{\Sigma}}_1 = \mathbf{I}_n \otimes \boldsymbol{\Sigma}$ in (1) yields directly to (3). Notice that, unlike for the probit example introduced in Section 2, in these cases $\bar{n}_0 = 0$, meaning that all the Gaussian responses are fully observed and there are no dichotomized outcomes. This translates into the fact that no contribution by the Gaussian cumulative distribution function is present in likelihoods (2)–(3), which, therefore, simplify to multivariate Gaussian densities.

2.2. Probit, Multivariate Probit and Multinomial Probit

As discussed in Section 2, the classical probit regression model $(y_i|\boldsymbol{\beta}) \sim \operatorname{Bern}[\Phi(\mathbf{x}_i^\top \boldsymbol{\beta})]$, independently for $i = 1, \dots, n$, induces likelihoods which can be readily reframed within representation (1). More specifically, recalling Durante (2019) and denoting with $\mathbf{1}_n$ the $(n \times 1)$ -dimensional vector of all ones, the probit likelihood can be expressed as

$$p(\mathbf{y}|\boldsymbol{\beta}) = \prod_{i=1}^n \Phi[(2y_i - 1)\mathbf{x}_i^\top \boldsymbol{\beta}] = \Phi_n(\operatorname{diag}(2\mathbf{y} - \mathbf{1}_n)\mathbf{X}\boldsymbol{\beta}; \mathbf{I}_n), \tag{4}$$

which is a special case of (1), with $\bar{n}_1 = 0$, $\bar{n}_0 = n$, $\bar{\mathbf{y}}_0 = \mathbf{0}$, $\bar{\mathbf{X}}_0 = \operatorname{diag}(2\mathbf{y} - \mathbf{1}_n)\mathbf{X}$ and $\bar{\boldsymbol{\Sigma}}_0 = \mathbf{I}_n$. Replacing $\bar{\boldsymbol{\Sigma}}_0 = \mathbf{I}_n$ with $\bar{\boldsymbol{\Sigma}}_0 = \sigma^2\mathbf{I}_n$ yields also probabilities of the form $\Phi(\mathbf{x}_i^\top \boldsymbol{\beta}; \sigma^2)$.

The above probit regression model further admits a number of routinely-used extensions which incorporate multivariate binary outcomes (Chib and Greenberg 1998) and also multinomial response data (Hausman and Wise 1978; Tutz 1991; Stern 1992). As previously mentioned, both cases have their roots in discrete choice models (e.g., Greene 2008), and can be reframed within (1). To clarify this result, let us first focus on multivariate probit models for the binary response vector $\mathbf{y}_i = (y_{i1}, \dots, y_{im})^\top \in \{0, 1\}^m$. As discussed in Chib and Greenberg (1998), these formulations can be interpreted as a dichotomized version of the regression model for multivariate Gaussian response data in Section 2.1. In fact, each \mathbf{y}_i is defined as $\mathbf{y}_i = [\mathbb{1}(z_{i1} > 0), \dots, \mathbb{1}(z_{im} > 0)]^\top = \mathbf{1}(\mathbf{z}_i > \mathbf{0})$, where $(\mathbf{z}_i = (z_{i1}, \dots, z_{im})^\top | \boldsymbol{\beta}) \sim N_m(\mathbf{X}_i\boldsymbol{\beta}, \boldsymbol{\Sigma})$, independently for every $i = 1, \dots, n$. This means that the contribution to the likelihood of each unit i is $p(\mathbf{y}_i|\boldsymbol{\beta}) = p(\mathbf{1}(\mathbf{z}_i > \mathbf{0})|\boldsymbol{\beta})$, which can be also written as $\Phi_m(\mathbf{B}_i\mathbf{X}_i\boldsymbol{\beta}; \mathbf{B}_i\boldsymbol{\Sigma}\mathbf{B}_i)$, following standard properties of multivariate Gaussian cumulative distribution functions, where $\mathbf{B}_i = \operatorname{diag}(2y_{i1} - 1, \dots, 2y_{im} - 1)$. As a result, the joint likelihood of multivariate probit is

$$p(\mathbf{y}|\boldsymbol{\beta}) = \prod_{i=1}^n \Phi_m(\mathbf{B}_i\mathbf{X}_i\boldsymbol{\beta}; \mathbf{B}_i\boldsymbol{\Sigma}\mathbf{B}_i) = \Phi_{n \cdot m}(\mathbf{B}\mathbf{X}\boldsymbol{\beta}; \mathbf{B}(\mathbf{I}_n \otimes \boldsymbol{\Sigma})\mathbf{B}), \tag{5}$$

where $\mathbf{X} = (\mathbf{X}_1^\top, \dots, \mathbf{X}_n^\top)^\top$, and \mathbf{B} denotes an $(n \cdot m) \times (n \cdot m)$ block-diagonal matrix with generic block $\mathbf{B}_{[i,i]} = \mathbf{B}_i$, for each $i = 1, \dots, n$. To reframe (5) within the general likelihood form

in (1) it suffices to set $\bar{n}_1 = 0$, $\bar{n}_0 = n \cdot m$, $\bar{y}_0 = \mathbf{0}$, $\bar{\mathbf{X}}_0 = \mathbf{B}\mathbf{X}$ and $\bar{\boldsymbol{\Sigma}}_0 = \mathbf{B}(\mathbf{I}_n \otimes \boldsymbol{\Sigma})\mathbf{B}$.

As discussed in Fasano and Durante (2022), a similar construction and derivations can be also considered for several multinomial probit models (Hausman and Wise 1978; Tutz 1991; Stern 1992). All these formulations express the probabilities of the L different categories $\{1, \dots, L\}$ via a discrete choice mechanism relying on correlated predictor-dependent Gaussian latent utilities which facilitate improved flexibility and avoid restrictive assumptions often found in multinomial logit, such as the independence of irrelevant alternatives (Hausman and Wise 1978). For instance, in the formulation by Stern (1992), each categorical response y_i is defined as $y_i = \operatorname{argmax}_l\{z_{il}, \dots, z_{iL}\}$, where $z_{il} = \mathbf{x}_i^\top \boldsymbol{\beta}_l + \varepsilon_{il}$ for $l = 1, \dots, L$, with $\boldsymbol{\varepsilon}_i = (\varepsilon_{i1}, \dots, \varepsilon_{iL})^\top \sim N_L(\mathbf{0}, \boldsymbol{\Sigma})$, and $\boldsymbol{\beta}_L = \mathbf{0}$ for identifiability purposes (Johndrow, Dunson, and Lum 2013). As a consequence of this construction and recalling Section 2.2 in Fasano and Durante (2022), it follows that $\operatorname{pr}(y_i = l|\boldsymbol{\beta}) = \operatorname{pr}(z_{il} > z_{ik}, \forall k \neq l)$, which can be also rewritten as $\operatorname{pr}(\varepsilon_{ik} - \varepsilon_{il} < \mathbf{x}_i^\top \boldsymbol{\beta}_l - \mathbf{x}_i^\top \boldsymbol{\beta}_k, \forall k \neq l)$. Therefore, let \mathbf{v}_l denote the L -dimensional vector with value 1 in position l and 0 elsewhere, for each $l = 1, \dots, L$, and define $\mathbf{x}_{il} = \bar{\mathbf{v}}_l \otimes \mathbf{x}_i$, where $\bar{\mathbf{v}}_l$ is the $(L - 1)$ -dimensional vector obtained by removing the l th entry in \mathbf{v}_l . Then, $\operatorname{pr}(y_i = l|\boldsymbol{\beta}) = \operatorname{pr}[(\mathbf{v}_k - \mathbf{v}_l)^\top \boldsymbol{\varepsilon}_i < (\mathbf{x}_{il} - \mathbf{x}_{ik})^\top \boldsymbol{\beta}, \forall k \neq l]$, where $\boldsymbol{\beta} = (\boldsymbol{\beta}_1^\top, \dots, \boldsymbol{\beta}_{L-1}^\top)^\top$. This expression can be also re-formulated in the more compact form $\operatorname{pr}(\mathbf{V}_{[-l]} \boldsymbol{\varepsilon}_i < \mathbf{X}_{i[-l]} \boldsymbol{\beta})$, where $\mathbf{V}_{[-l]}$ and $\mathbf{X}_{i[-l]}$ correspond to suitable matrices whose rows are obtained by stacking the vectors $(\mathbf{v}_k - \mathbf{v}_l)^\top$ and $(\mathbf{x}_{il} - \mathbf{x}_{ik})^\top$, respectively, for every $k \neq l$, and $<$ is intended elementwise. Therefore, leveraging the standard properties of multivariate Gaussians, it follows that $\operatorname{pr}(y_i = l|\boldsymbol{\beta}) = \operatorname{pr}(\mathbf{V}_{[-l]} \boldsymbol{\varepsilon}_i < \mathbf{X}_{i[-l]} \boldsymbol{\beta}) = \Phi_{L-1}(\mathbf{X}_{i[-l]} \boldsymbol{\beta}; \mathbf{V}_{[-l]} \boldsymbol{\Sigma} \mathbf{V}_{[-l]}^\top)$, for every $l = 1, \dots, L$. This result yields a joint likelihood for the categorical responses $\mathbf{y} = (y_1, \dots, y_n)^\top$ which can be written as

$$p(\mathbf{y}|\boldsymbol{\beta}) = \prod_{i=1}^n \Phi_{L-1}(\mathbf{X}_{i[-y_i]} \boldsymbol{\beta}; \mathbf{V}_{[-y_i]} \boldsymbol{\Sigma} \mathbf{V}_{[-y_i]}^\top) = \Phi_{n \cdot (L-1)}(\mathbf{X}\boldsymbol{\beta}; \mathbf{V}(\mathbf{I}_n \otimes \boldsymbol{\Sigma})\mathbf{V}^\top), \tag{6}$$

where $\mathbf{X} = (\mathbf{X}_{1[-y_1]}^\top, \dots, \mathbf{X}_{n[-y_n]}^\top)^\top$, and \mathbf{V} is a block-diagonal matrix with generic block $\mathbf{V}_{[i,l]} = \mathbf{V}_{[-y_i]}$, for each $i = 1, \dots, n$. Setting $\bar{n}_1 = 0$, $\bar{n}_0 = n \cdot (L - 1)$, $\bar{\mathbf{y}}_0 = \mathbf{0}$, $\bar{\mathbf{X}}_0 = \mathbf{X}$ and $\bar{\boldsymbol{\Sigma}}_0 = \mathbf{V}(\mathbf{I}_n \otimes \boldsymbol{\Sigma})\mathbf{V}^\top$ in (1) leads to (6). Hence, the multinomial probit model by Stern (1992) is again a special case of the general form in (1). As shown in Sections 2.1 and 2.3 of Fasano and Durante (2022), also the alternative formulations proposed by Hausman and Wise (1978) and Tutz (1991) induce likelihoods which can be expressed as cumulative distribution functions of multivariate Gaussians evaluated at a suitable linear combination of the coefficients' vector $\boldsymbol{\beta}$; see Propositions 1 and 3 in Fasano and Durante (2022). This means that also such models can be easily recasted within the general form in (1) with $\bar{n}_1 = 0$, $\bar{\mathbf{y}}_0 = \mathbf{0}$, and suitably defined \bar{n}_0 , $\bar{\mathbf{X}}_0$ and $\bar{\boldsymbol{\Sigma}}_0$.

2.3. Tobit Regression

Recalling Section 2, the classical tobit model (Tobin 1958) characterizes the intermediate situation in which response data are fully observed only if exceeding a certain threshold, often set to 0. This means that $y_i = z_i \mathbb{1}(z_i > 0)$, with $(z_i|\boldsymbol{\beta}) \sim N(\mathbf{x}_i^\top \boldsymbol{\beta}, \sigma^2)$,

independently for $i = 1, \dots, n$. Such a formulation yields the joint likelihood

$$p(\mathbf{y}|\boldsymbol{\beta}) = \prod_{i:y_i>0} \phi(y_i - \mathbf{x}_i^\top \boldsymbol{\beta}; \sigma^2) \prod_{i:y_i=0} \Phi(-\mathbf{x}_i^\top \boldsymbol{\beta}; \sigma^2) = \phi_{n_1}(\mathbf{y}_1 - \mathbf{X}_1 \boldsymbol{\beta}; \sigma^2 \mathbf{I}_{n_1}) \Phi_{n_0}(-\mathbf{X}_0 \boldsymbol{\beta}; \sigma^2 \mathbf{I}_{n_0}), \tag{7}$$

where n_1 and n_0 denote the number of fully observed and censored units, respectively, whereas \mathbf{y}_1 , \mathbf{X}_1 , and \mathbf{X}_0 are the response vector and design matrices associated with these two subsets of units. This likelihood can be again expressed as a special example of (1) by letting $\bar{n}_1 = n_1$, $\bar{n}_0 = n_0$, $\bar{\mathbf{y}}_1 = \mathbf{y}_1$, $\bar{\mathbf{y}}_0 = \mathbf{0}$, $\bar{\mathbf{X}}_1 = \mathbf{X}_1$, $\bar{\mathbf{X}}_0 = -\mathbf{X}_0$, $\bar{\boldsymbol{\Sigma}}_1 = \sigma^2 \mathbf{I}_{n_1}$, and $\bar{\boldsymbol{\Sigma}}_0 = \sigma^2 \mathbf{I}_{n_0}$.

The above result also holds for several subsequent extensions of the original tobit model (Tobin 1958), which include more elaborated censoring mechanisms, possibly relying on multivariate Gaussian utilities. Such generalizations, often known in the literature as type II, III, IV, and V tobit models, are carefully discussed in Amemiya (1984) and all induce likelihoods which can be written as the product of Gaussian densities and cumulative distribution functions evaluated at suitable linear combinations of the coefficients $\boldsymbol{\beta}$. This common structure allows again to readily express such extensions as special cases of the general form in (1). Inclusion of multivariate versions is also straightforward under a similar reasoning considered in (3) and (5).

2.4. Extensions to Skewed, Nonlinear and Dynamic Models

As discussed in Section S1 of the supplementary materials, although the models discussed in Sections 2.1–2.3 cover the most widely-implemented formulations in the literature, several additional extensions of these representations to skewed (e.g., Chen, Dey, and Shao 1999; Sahu, Dey, and Branco 2003; Bazán, Bolfarine, and Branco 2010; Hutton and Stanghellini 2011), nonlinear (e.g., Kuss, Rasmussen, and Herbrich 2005; De Oliveira 2005; Nickisch and Rasmussen 2008; Riihimäki, Jylänki, and Vehtari 2014; Cao, Durante, and Genton 2022; Benavoli, Azzimonti, and Piga 2020), dynamic (e.g., Manrique and Shephard 1998; Andrieu and Doucet 2002; Naveau, Genton, and Shen 2005; Chib and Jeliazkov 2006; Soyer and Sung 2013; Fasano et al. 2021) and other contexts admit likelihood forms as in (1).

3. Conjugacy via Unified Skew-Normals

Sections 3.1–3.2 unify Bayesian inference for the whole family of models in Section 2 by proving that the likelihood in (1) admits as conjugate priors the class of unified skew-normal (SUN) distributions (Arellano-Valle and Azzalini 2006). Crucially, these variables include as special cases the commonly-assumed Gaussian priors for $\boldsymbol{\beta}$ in models (2)–(7), while extending such distributions in several directions. Hence, this review not only unifies and extends a broad class of models within a single likelihood representation, but also enlarges the class of prior distributions which admit closed-form posteriors that facilitate Bayesian inference.

3.1. Unified Skew-Normal Prior

Classical Bayesian implementations of the models in Sections 2.1–2.4 often assume multivariate Gaussian priors for $\boldsymbol{\beta}$,

which are natural choices in Bayesian regression and, under the models presented in Sections 2.1–2.4, are further motivated by the Gaussian form of the underlying latent utilities (e.g., Chib 1992; Albert and Chib 1993; McCulloch and Rossi 1994; Nobile 1998; Chib and Greenberg 1998; McCulloch, Polson, and Rossi 2000; Albert and Chib 2001; Imai and Van Dyk 2005; Kuss, Rasmussen, and Herbrich 2005; Holmes and Held 2006; Riihimäki, Jylänki, and Vehtari 2014; Chopin and Ridgway 2017). Interestingly, these priors are special cases of more general distributions which induce asymmetric shapes in multivariate Gaussians by modifying the density of such variables through a skewness-inducing mechanism driven by the cumulative distribution function of another Gaussian. Key examples include multivariate skew-normals (Azzalini and Dalla Valle 1996; Azzalini and Capitanio 1999), extended multivariate skew-normals (Arnold and Beaver 2000; Arnold et al. 2002) and closed skew-normals (González-Farias, Dominguez-Molina, and Gupta 2004; Gupta, González-Farias, and Dominguez-Molina 2004), which have all been subsequently unified by Arellano-Valle and Azzalini (2006) in a single general class, namely the unified skew-normal (SUN) distribution. Recalling Arellano-Valle and Azzalini (2006), the vector $\beta \in \mathbb{R}^{\bar{p}}$ has $\text{SUN}_{\bar{p},\bar{n}}(\xi, \Omega, \Delta, \gamma, \Gamma)$ prior if its density $p(\beta)$ is equal to

$$\phi_{\bar{p}}(\beta - \xi; \Omega) \frac{\Phi_{\bar{n}}(\gamma + \Delta^T \bar{\Omega}^{-1} \omega^{-1}(\beta - \xi); \Gamma - \Delta^T \bar{\Omega}^{-1} \Delta)}{\Phi_{\bar{n}}(\gamma; \Gamma)}, \tag{8}$$

with $\bar{\Omega}$ denoting the $\bar{p} \times \bar{p}$ correlation matrix associated with the covariance matrix Ω which, in turn, can be expressed as $\Omega = \omega \bar{\Omega} \omega$, where $\omega = (\Omega \odot \mathbf{I}_{\bar{p}})^{1/2}$, and \odot refers to the element-wise Hadamard product. According to (8), skewness is induced in $\phi_{\bar{p}}(\beta - \xi; \Omega)$ by multiplying such a density with the cumulative distribution function of a $N_{\bar{n}}(\mathbf{0}, \Gamma - \Delta^T \bar{\Omega}^{-1} \Delta)$, evaluated at $\gamma + \Delta^T \bar{\Omega}^{-1} \omega^{-1}(\beta - \xi)$, whereas $\Phi_{\bar{n}}(\gamma; \Gamma)$ corresponds to the normalizing constant. Notice that when all the entries in the $\bar{p} \times \bar{n}$ skewness matrix Δ are 0, the numerator in (8) reduces to $\Phi_{\bar{n}}(\gamma; \Gamma)$, thereby allowing to obtain the classical Gaussian prior density $\phi_{\bar{p}}(\beta - \xi; \Omega)$ as a special case of (8). The quantities \bar{p} and \bar{n} denote instead the dimensions of the density and the cumulative distribution function, respectively. Within the general class of formulations discussed in Sections 2.1–2.4, \bar{p} refers to the dimension of β and, hence, can vary depending on the model considered. While in most cases \bar{p} is equal to the number of predictors p , under specific constructions the two dimensions might differ. For instance, in the multinomial probit model in (6), \bar{p} coincides with $p \cdot (L - 1)$. Conversely, \bar{n} defines the dimension of the multivariate Gaussian cumulative distribution function responsible for the skewness-inducing mechanism in the prior density. For example, setting $\bar{n} = 0$ yields the classical Gaussian prior for β , whereas assuming independent skew-normals for each $\beta_j, j = 1, \dots, \bar{p}$, would imply $\bar{n} = \bar{p}$.

Recalling Arellano-Valle and Azzalini (2006), the above SUN distribution also admits a generative construction which further clarifies the role of the parameters $\xi, \Omega, \Delta, \gamma$, and Γ , and provides key intuitions on the conjugacy properties of SUN priors under likelihood (1). In particular, let $\tilde{\beta} \in \mathbb{R}^{\bar{p}}$ and $\tilde{z} \in \mathbb{R}^{\bar{n}}$ denote two vectors jointly distributed as a $N_{\bar{p}+\bar{n}}(\mathbf{0}, \Omega^*)$, where

Ω^* is a $(\bar{p}+\bar{n}) \times (\bar{p}+\bar{n})$ correlation matrix with blocks $\Omega_{[11]}^* = \bar{\Omega}$, $\Omega_{[22]}^* = \Gamma$ and $\Omega_{[21]}^* = \Omega_{[12]}^{*\top} = \Delta^T$, then $\tilde{\beta} = (\tilde{\beta} | \tilde{z} > \gamma > \mathbf{0})$ is distributed as a $\text{SUN}_{\bar{p},\bar{n}}(\mathbf{0}, \bar{\Omega}, \Delta, \gamma, \Gamma)$, while $\beta = \xi + \omega \tilde{\beta} \sim \text{SUN}_{\bar{p},\bar{n}}(\xi, \Omega, \Delta, \gamma, \Gamma)$ with density as in (8). Consistent with this generative representation, the parameters ξ and ω control the location and the scale of the prior, whereas $\bar{\Omega}, \Gamma$, and Δ regulate the dependence within $\tilde{\beta}, \tilde{z}$ and between these two random vectors, respectively. The term γ denotes instead the truncation threshold in the conditioning mechanism. Besides clarifying the role of the prior parameters, this representation offers intuitions on the SUN conjugacy formalized in Section 3.2. In fact, according to such a construction, SUNs arise as conditional distributions in a generative mechanism that relies on partially-observed Gaussian latent variables \tilde{z} . This has direct connections with the posterior distribution for β under the models in Section 2, which is also defined, through Bayes rule, via a conditioning operation relying on partially or fully observed Gaussian utilities.

The above discussion further suggests that different forms of prior information can be included via (8). Since multivariate normal distributions are special cases of SUNs, all non-informative, weakly informative and informative priors relying on Gaussians (e.g., Zellner 1986; Gelman et al. 2008; Chopin and Ridgway 2017) can be employed by letting $\bar{n} = 0$, and suitably specifying ξ and Ω . The possibility to include skewness further allows the incorporation of additional prior information by letting $\bar{n} > 0$ and choosing appropriate values for the parameters Δ, γ , and Γ , keeping in mind the corresponding role in the generative process that leads to the prior density in (8). Key examples of priors of potential interest in this context, which belong to the SUN family, are univariate skew-normals (Azzalini 1985) for each coefficient $\beta_j, j = 1, \dots, \bar{p}$, or multivariate skew-normals (Azzalini and Dalla Valle 1996; Azzalini and Capitanio 1999), extended multivariate skew-normals (Arnold and Beaver 2000; Arnold et al. 2002) and closed skew-normals (González-Farias, Dominguez-Molina, and Gupta 2004; Gupta, González-Farias, and Dominguez-Molina 2004) for the entire vector β . Among such options, independent univariate skew-normals are a convenient choice that easily allows to elicit skewness information for each $\beta_j, j = 1, \dots, \bar{p}$ via a single and interpretable parameter.

As clarified in Section 3.2, the SUNs' properties are also beneficial for posterior inference. Recalling Arellano-Valle and Azzalini (2006), Azzalini and Bacchieri (2010), Azzalini and Capitanio (2013), and Arellano-Valle and Azzalini (2021) SUNs have a number of properties in common with multivariate Gaussians. These include closure under marginalization, linear combinations and conditioning, along with the availability of closed-form expressions for the moment generating function, and additive representations via linear combinations of multivariate Gaussians and truncated normals. Due to the SUN conjugacy proved in Theorem 1, all these properties can facilitate point estimation, uncertainty quantification, model selection and prediction under the SUN posterior associated with the general likelihood in (1) which encompasses the models in Sections 2.1–2.4. This provides important advancements for a broad class of models, under a similarly wide family of priors beyond Gaussians.

3.2. Unified Skew-Normal Posterior and its Properties

Theorem 1 unifies and extends recent model-specific derivations by proving SUN conjugacy for any statistical model whose likelihood can be expressed as in (1). The proof of **Theorem 1** combines original results on SUN conjugacy in probit models (Durante 2019, Corollary 4) with **Lemma 1**, which shows that SUN priors are also conjugate to Gaussian linear regression; see Section S2 in the supplementary materials for detailed proofs of **Lemma 1** and **Theorem 1**.

Lemma 1. Let $p(\bar{y}_1|\boldsymbol{\beta}) = \phi_{\bar{n}_1}(\bar{y}_1 - \bar{\mathbf{X}}_1\boldsymbol{\beta}; \bar{\boldsymbol{\Sigma}}_1)$ and assume that $\boldsymbol{\beta}$ has $\text{SUN}_{\bar{p},\bar{n}}(\boldsymbol{\xi}, \boldsymbol{\Omega}, \boldsymbol{\Delta}, \boldsymbol{\gamma}, \boldsymbol{\Gamma})$ prior with density $p(\boldsymbol{\beta})$ as in (8). Then, $(\boldsymbol{\beta}|\bar{y}_1) \sim \text{SUN}_{\bar{p},\bar{n}}(\boldsymbol{\xi}_1, \boldsymbol{\Omega}_1, \boldsymbol{\Delta}_1, \boldsymbol{\gamma}_1, \boldsymbol{\Gamma}_1)$ where

$$\begin{aligned}\boldsymbol{\xi}_1 &= (\boldsymbol{\Omega}^{-1} + \bar{\mathbf{X}}_1^\top \bar{\boldsymbol{\Sigma}}_1^{-1} \bar{\mathbf{X}}_1)^{-1} (\boldsymbol{\Omega}^{-1} \boldsymbol{\xi} + \bar{\mathbf{X}}_1^\top \bar{\boldsymbol{\Sigma}}_1^{-1} \bar{y}_1), \\ \boldsymbol{\Omega}_1 &= (\boldsymbol{\Omega}^{-1} + \bar{\mathbf{X}}_1^\top \bar{\boldsymbol{\Sigma}}_1^{-1} \bar{\mathbf{X}}_1)^{-1} = \boldsymbol{\omega}_1 \bar{\boldsymbol{\Omega}}_1 \boldsymbol{\omega}_1, \\ \boldsymbol{\Delta}_1 &= \bar{\boldsymbol{\Omega}}_1 \boldsymbol{\omega}_1 \boldsymbol{\omega}_1^{-1} \bar{\boldsymbol{\Omega}}_1^{-1} \boldsymbol{\Delta} \mathbf{s}_1^{-1}, \\ \boldsymbol{\gamma}_1 &= \mathbf{s}_1^{-1} [\boldsymbol{\gamma} + \boldsymbol{\Delta}^\top \bar{\boldsymbol{\Omega}}_1^{-1} \boldsymbol{\omega}_1^{-1} (\boldsymbol{\xi}_1 - \boldsymbol{\xi})], \\ \boldsymbol{\Gamma}_1 &= \mathbf{s}_1^{-1} [\boldsymbol{\Gamma} + \boldsymbol{\Delta}^\top (\bar{\boldsymbol{\Omega}}_1^{-1} \boldsymbol{\omega}_1^{-1} \bar{\boldsymbol{\Omega}}_1 \boldsymbol{\omega}_1^{-1} \bar{\boldsymbol{\Omega}}_1^{-1} - \bar{\boldsymbol{\Omega}}_1^{-1}) \boldsymbol{\Delta}] \mathbf{s}_1^{-1},\end{aligned}$$

with $\mathbf{s}_1 = ((\boldsymbol{\Gamma} + \boldsymbol{\Delta}^\top (\bar{\boldsymbol{\Omega}}_1^{-1} \boldsymbol{\omega}_1^{-1} \bar{\boldsymbol{\Omega}}_1 \boldsymbol{\omega}_1^{-1} \bar{\boldsymbol{\Omega}}_1^{-1} - \bar{\boldsymbol{\Omega}}_1^{-1}) \boldsymbol{\Delta}) \odot \mathbf{I}_{\bar{n}})^{1/2}$.

Note that in **Lemma 1** the rescaling operated by \mathbf{s}_1 is required to ensure that the matrix $\boldsymbol{\Omega}_1^*$ with blocks $\boldsymbol{\Omega}_{1[11]}^* = \bar{\boldsymbol{\Omega}}_1$, $\boldsymbol{\Omega}_{1[22]}^* = \boldsymbol{\Gamma}_1$ and $\boldsymbol{\Omega}_{1[21]}^* = \boldsymbol{\Omega}_{1[12]}^{*\top} = \boldsymbol{\Delta}_1^\top$ is a correlation matrix, as in the original formulation (Arellano-Valle and Azzalini 2006). Although this constraint is useful to avoid identifiability issues in frequentist contexts, such problems are less of a concern in our Bayesian setting since the parameters of the SUN posterior are known functions of the observed data and of the pre-specified prior hyperparameters, and thus do not need to be estimated. However, maintaining this constraint is still useful to inherit results of the original SUN and to avoid identifiability issues in prior elicitation.

Leveraging **Lemma 1** above, and adapting Corollary 4 in Durante (2019), it is now possible to state the general SUN conjugacy result in **Theorem 1**.

Theorem 1. Consider the likelihood

$$p(\mathbf{y}|\boldsymbol{\beta}) \propto \phi_{\bar{n}_1}(\bar{y}_1 - \bar{\mathbf{X}}_1\boldsymbol{\beta}; \bar{\boldsymbol{\Sigma}}_1) \Phi_{\bar{n}_0}(\bar{y}_0 + \bar{\mathbf{X}}_0\boldsymbol{\beta}; \bar{\boldsymbol{\Sigma}}_0)$$

defined in (1), and assume that $\boldsymbol{\beta}$ has a $\text{SUN}_{\bar{p},\bar{n}}(\boldsymbol{\xi}, \boldsymbol{\Omega}, \boldsymbol{\Delta}, \boldsymbol{\gamma}, \boldsymbol{\Gamma})$ prior with density as in (8). Then

$$(\boldsymbol{\beta}|\mathbf{y}) \sim \text{SUN}_{\bar{p},\bar{n}+\bar{n}_0}(\boldsymbol{\xi}_{\text{post}}, \boldsymbol{\Omega}_{\text{post}}, \boldsymbol{\Delta}_{\text{post}}, \boldsymbol{\gamma}_{\text{post}}, \boldsymbol{\Gamma}_{\text{post}})$$

with posterior parameters

$$\begin{aligned}\boldsymbol{\xi}_{\text{post}} &= \boldsymbol{\xi}_1, & \boldsymbol{\Omega}_{\text{post}} &= \boldsymbol{\Omega}_1, & \boldsymbol{\Delta}_{\text{post}} &= (\boldsymbol{\Delta}_1, \bar{\boldsymbol{\Omega}}_1 \boldsymbol{\omega}_1 \bar{\mathbf{X}}_0^\top \mathbf{s}_0^{-1}), \\ \boldsymbol{\gamma}_{\text{post}} &= (\boldsymbol{\gamma}_1^\top, (\bar{y}_0 + \bar{\mathbf{X}}_0 \boldsymbol{\xi}_1)^\top \mathbf{s}_0^{-1})^\top,\end{aligned}$$

and $\boldsymbol{\Gamma}_{\text{post}}$ denoting a full-rank $(\bar{n} + \bar{n}_0) \times (\bar{n} + \bar{n}_0)$ correlation matrix with blocks

$$\begin{aligned}\boldsymbol{\Gamma}_{\text{post}[11]} &= \boldsymbol{\Gamma}_1, & \boldsymbol{\Gamma}_{\text{post}[22]} &= \mathbf{s}_0^{-1} (\bar{\mathbf{X}}_0 \boldsymbol{\Omega}_1 \bar{\mathbf{X}}_0^\top + \bar{\boldsymbol{\Sigma}}_0) \mathbf{s}_0^{-1}, \\ \boldsymbol{\Gamma}_{\text{post}[21]} &= \boldsymbol{\Gamma}_{\text{post}[12]}^\top = \mathbf{s}_0^{-1} \bar{\mathbf{X}}_0 \boldsymbol{\omega}_1 \boldsymbol{\Delta}_1,\end{aligned}$$

where $\mathbf{s}_0 = [(\bar{\mathbf{X}}_0 \boldsymbol{\Omega}_1 \bar{\mathbf{X}}_0^\top + \bar{\boldsymbol{\Sigma}}_0) \odot \mathbf{I}_{\bar{n}_0}]^{1/2}$, while $\boldsymbol{\xi}_1$, $\boldsymbol{\Omega}_1$, $\boldsymbol{\Delta}_1$, $\boldsymbol{\gamma}_1$, and $\boldsymbol{\Gamma}_1$ are defined as in **Lemma 1**.

Theorem 1 encompasses all available conjugacy results for SUN distributions under specific models within the broader family analyzed, while extending such findings to other key formulations. For example, setting $\bar{p} = p$, $\bar{n}_1 = 0$, $\bar{n}_0 = n$, $\bar{y}_0 = \mathbf{0}$, $\bar{\mathbf{X}}_0 = \text{diag}(2\mathbf{y} - \mathbf{1}_n)\mathbf{X}$ and $\bar{\boldsymbol{\Sigma}}_0 = \mathbf{I}_n$ as in (4), and substituting these quantities within the expressions in **Theorem 1**, would yield a $\text{SUN}_{p,\bar{n}+n}(\boldsymbol{\xi}_{\text{post}}, \boldsymbol{\Omega}_{\text{post}}, \boldsymbol{\Delta}_{\text{post}}, \boldsymbol{\gamma}_{\text{post}}, \boldsymbol{\Gamma}_{\text{post}})$ posterior with parameters as in Corollary 4 by Durante (2019). **Theorem 1** in Durante (2019) is instead recovered under the additional constraint $\bar{n} = 0$, which implies a Gaussian prior. Note that when $\bar{n}_1 = 0$ the associated quantities \bar{y}_1 , $\bar{\mathbf{X}}_1$, and $\bar{\boldsymbol{\Sigma}}_1$ are not defined and simply need to be removed from the formulas in **Theorem 1**. The same reasoning holds for \bar{y}_0 , $\bar{\mathbf{X}}_0$, and $\bar{\boldsymbol{\Sigma}}_0$ when $\bar{n}_0 = 0$, and for $\boldsymbol{\Delta}$, $\boldsymbol{\gamma}$ and $\boldsymbol{\Gamma}$ if $\bar{n} = 0$. For instance, setting $\bar{n}_0 = 0$ in **Theorem 1** leads to **Lemma 1**. Similarly, the SUN conjugacy results for multinomial probit (Fasano and Durante 2022), dynamic multivariate probit (Fasano et al. 2021), Gaussian processes (Cao, Durante, and Genton 2022), and skewed Gaussian processes under linear models, affine probit and combinations of these two formulations (Benavoli, Azzimonti, and Piga 2020, 2021) can be readily obtained from **Theorem 1** under the settings in Sections 2.1–2.4 for the quantities defining the likelihood in (1). Interestingly, also results outside of the regression context, such as those proved by Canale, Pagui, and Scarpa (2016) for multivariate skew-normal likelihoods with Gaussian or skew-normal priors on the shape parameter, can be recasted within **Theorem 1**. Besides encompassing already available findings, **Theorem 1** provides novel conjugacy results also in previously-unexplored settings, such as in tobit regression and in models relying on skewed utilities.

As discussed in Section 3.1, the availability of a SUN posterior in **Theorem 1** facilitates Bayesian inference for the whole class of models in Sections 2.1–2.4, by leveraging known properties of SUNs (Azzalini and Capitanio 2013; Arellano-Valle and Azzalini 2021). For instance, recalling Arellano-Valle and Azzalini (2006), the moment generating function of the posterior is

$$M(\mathbf{t}) = e^{\boldsymbol{\xi}_{\text{post}}^\top \mathbf{t} + 0.5 \mathbf{t}^\top \boldsymbol{\Omega}_{\text{post}} \mathbf{t}} \frac{\Phi_{\bar{n}+\bar{n}_0}(\boldsymbol{\gamma}_{\text{post}} + \boldsymbol{\Delta}_{\text{post}}^\top \boldsymbol{\omega}_{\text{post}} \mathbf{t}; \boldsymbol{\Gamma}_{\text{post}})}{\Phi_{\bar{n}+\bar{n}_0}(\boldsymbol{\gamma}_{\text{post}}; \boldsymbol{\Gamma}_{\text{post}})}, \quad (9)$$

for $\mathbf{t} \in \mathbb{R}^{\bar{p}}$, and, therefore, closed-form expressions for relevant moments can be obtained from (9). In particular, applying the derivations in Azzalini and Bacchieri (2010) and Arellano-Valle and Azzalini (2021) to the SUN posterior in **Theorem 1**, yields the following expressions for $\mathbb{E}(\boldsymbol{\beta}|\mathbf{y})$ and $\text{var}(\boldsymbol{\beta}|\mathbf{y})$

$$\begin{aligned}\mathbb{E}(\boldsymbol{\beta}|\mathbf{y}) &= \boldsymbol{\xi}_{\text{post}} + \boldsymbol{\omega}_{\text{post}} \boldsymbol{\Delta}_{\text{post}} \boldsymbol{\psi}, \\ \text{var}(\boldsymbol{\beta}|\mathbf{y}) &= \boldsymbol{\Omega}_{\text{post}} + \boldsymbol{\omega}_{\text{post}} \boldsymbol{\Delta}_{\text{post}} (\boldsymbol{\Psi} - \boldsymbol{\psi} \boldsymbol{\psi}^\top) \boldsymbol{\Delta}_{\text{post}}^\top \boldsymbol{\omega}_{\text{post}},\end{aligned} \quad (10)$$

where $\boldsymbol{\psi}$ is a vector of dimensions $(\bar{n} + \bar{n}_0) \times 1$, with entries $\psi_i = \phi(\boldsymbol{\gamma}_{\text{post},i}) \Phi_{\bar{n}+\bar{n}_0-1}(\boldsymbol{\gamma}_{\text{post},-i} - \boldsymbol{\Gamma}_{\text{post},-i} \boldsymbol{\gamma}_{\text{post},i}; \boldsymbol{\Gamma}_{\text{post},-i,-i} - \boldsymbol{\Gamma}_{\text{post},-i} \boldsymbol{\Gamma}_{\text{post},-i}^\top) / \Phi_{\bar{n}+\bar{n}_0}(\boldsymbol{\gamma}_{\text{post}}; \boldsymbol{\Gamma}_{\text{post}})$ for $i = 1, \dots, \bar{n} + \bar{n}_0$, where $\boldsymbol{\gamma}_{\text{post},i}$ and $\boldsymbol{\gamma}_{\text{post},-i}$ correspond to the i th element of $\boldsymbol{\gamma}_{\text{post}}$ and the $(\bar{n} + \bar{n}_0 - 1) \times 1$ vector obtained by removing entry i in $\boldsymbol{\gamma}_{\text{post}}$, respectively, whereas $\boldsymbol{\Gamma}_{\text{post},-i}$ and $\boldsymbol{\Gamma}_{\text{post},-i,-i}$ are the i th column of $\boldsymbol{\Gamma}_{\text{post}}$ without entry i and the sub-matrix obtained by removing the i th row and column from $\boldsymbol{\Gamma}_{\text{post}}$, respectively. Analogously, $\boldsymbol{\Psi}$ is a $(\bar{n} + \bar{n}_0) \times (\bar{n} + \bar{n}_0)$ symmetric matrix involving the second-order derivatives of the cumulative distribution function

term in (9); refer to Arellano-Valle and Azzalini (2021) for the exact expression of Ψ and of higher-order moments of the SUN. These quantities can be also estimated via Monte Carlo since

$$(\beta|\mathbf{y}) \stackrel{d}{=} \xi_{\text{post}} + \omega_{\text{post}}(\mathbf{U}_0 + \Delta_{\text{post}}\Gamma_{\text{post}}^{-1}\mathbf{U}_1), \tag{11}$$

with $\stackrel{d}{=}$ meaning equality in distribution, and

$$\begin{aligned} \mathbf{U}_0 &\sim N_{\bar{p}}(\mathbf{0}, \bar{\Omega}_{\text{post}} - \Delta_{\text{post}}\Gamma_{\text{post}}^{-1}\Delta_{\text{post}}^{\top}), \\ \mathbf{U}_1 &\sim \text{TN}_{\bar{n}+\bar{n}_0}(-\boldsymbol{\gamma}_{\text{post}}; \mathbf{0}, \Gamma_{\text{post}}), \end{aligned}$$

where $\text{TN}_{\bar{n}+\bar{n}_0}(-\boldsymbol{\gamma}_{\text{post}}; \mathbf{0}, \Gamma_{\text{post}})$ denotes an $(\bar{n} + \bar{n}_0)$ -variate Gaussian having mean $\mathbf{0}$, covariance matrix Γ_{post} and truncation below $-\boldsymbol{\gamma}_{\text{post}}$. This additive construction has been first derived in Arellano-Valle and Azzalini (2006) and allows to generate independent and identically distributed values from the exact posterior via linear combinations of samples from \bar{p} -variate Gaussians and $(\bar{n} + \bar{n}_0)$ -variate truncated normals, thus overcoming convergence and mixing issues of MCMC methods; see Section 4.

Uncertainty quantification and calculation of credible intervals is instead facilitated by the availability of a closed-form expression for the SUN cumulative distribution function. Adapting Azzalini and Bacchieri (2010) and Arellano-Valle and Azzalini (2021), this is

$$\text{pr}(\beta \leq \mathbf{b}|\mathbf{y}) = \frac{\Phi_{\bar{p}+(\bar{n}+\bar{n}_0)}([\mathbf{b} - \xi_{\text{post}}]^{\top} \omega_{\text{post}}^{-1} \boldsymbol{\gamma}_{\text{post}}^{\top}]^{\top}; \tilde{\Omega}_{\text{post}})}{\Phi_{\bar{n}+\bar{n}_0}(\boldsymbol{\gamma}_{\text{post}}; \Gamma_{\text{post}})}, \tag{12}$$

for every $\mathbf{b} \in \mathbb{R}^{\bar{p}}$, where $\tilde{\Omega}_{\text{post}}$ denotes a matrix with blocks $\tilde{\Omega}_{\text{post}\{11\}} = \bar{\Omega}_{\text{post}}$, $\tilde{\Omega}_{\text{post}\{22\}} = \Gamma_{\text{post}}$ and $\tilde{\Omega}_{\text{post}\{21\}} = \tilde{\Omega}_{\text{post}\{12\}}^{\top} = -\Delta_{\text{post}}^{\top}$.

Extending the results of Durante (2019), Fasano and Durante (2022) and Benavoli, Azzimonti, and Piga (2021) to the more general setting under consideration, it is also possible to obtain the marginal likelihood as follows

$$p(\mathbf{y}) = c \cdot \phi_{\bar{n}_1}(\bar{\mathbf{y}}_1 - \bar{\mathbf{X}}_1 \xi; \bar{\Sigma}_1 + \bar{\mathbf{X}}_1 \Omega \bar{\mathbf{X}}_1^{\top}) \frac{\Phi_{\bar{n}+\bar{n}_0}(\boldsymbol{\gamma}_{\text{post}}; \Gamma_{\text{post}})}{\Phi_{\bar{n}}(\boldsymbol{\gamma}; \Gamma)}, \tag{13}$$

where $c = 1$ under all the routinely-implemented models in Section 2 which rely on Gaussian utilities, namely (2)–(7), whereas for those formulations based on skewed utilities, for example, (S.1)–(S.3), the constant c is a known value. Albeit the primary interest is inference on β , the marginal likelihood in (13) allows to obtain empirical Bayes estimates also for the other quantities in likelihood (1), such as the parameters of the covariance matrices $\bar{\Sigma}_1$ and $\bar{\Sigma}_0$, via numerical maximization; see Section 6 for further discussion on estimation of $\bar{\Sigma}_1$ and $\bar{\Sigma}_0$. In addition, (13) facilitates direct calculation of Bayes factors for model selection and evaluation of predictive probabilities. This second objective can be readily accomplished noting that the predictive probability $p(\mathbf{y}_{\text{new}}|\mathbf{y})$ for a new vector of observations \mathbf{y}_{new} from model (1) is equal to the ratio $p(\mathbf{y}_{\text{new}}, \mathbf{y})/p(\mathbf{y})$ of the two associated marginal likelihoods. Therefore, focusing for simplicity on the case $c = 1$, which covers the most widely-used models in Section 2, direct application of (13)

leads to

$$p(\mathbf{y}_{\text{new}}|\mathbf{y}) = \frac{\phi_{\bar{n}_1+\bar{n}_{1\text{new}}}(\bar{\mathbf{y}}_{1\text{pred}} - \bar{\mathbf{X}}_{1\text{pred}} \xi; \bar{\Sigma}_{1\text{pred}} + \bar{\mathbf{X}}_{1\text{pred}} \Omega \bar{\mathbf{X}}_{1\text{pred}}^{\top})}{\phi_{\bar{n}_1}(\bar{\mathbf{y}}_1 - \bar{\mathbf{X}}_1 \xi; \bar{\Sigma}_1 + \bar{\mathbf{X}}_1 \Omega \bar{\mathbf{X}}_1^{\top})} \cdot \frac{\Phi_{\bar{n}+\bar{n}_0+\bar{n}_{0\text{new}}}(\boldsymbol{\gamma}_{\text{pred}}; \Gamma_{\text{pred}})}{\Phi_{\bar{n}+\bar{n}_0}(\boldsymbol{\gamma}_{\text{post}}; \Gamma_{\text{post}})}, \tag{14}$$

where $\bar{n}_{1\text{new}}$ and $\bar{n}_{0\text{new}}$ are the dimensions of the two vectors $\bar{\mathbf{y}}_{1\text{new}}$ and $\bar{\mathbf{y}}_{0\text{new}}$ associated with \mathbf{y}_{new} . Similarly, $\bar{\mathbf{y}}_{1\text{pred}} = (\bar{\mathbf{y}}_1^{\top}, \bar{\mathbf{y}}_{1\text{new}}^{\top})^{\top}$, $\bar{\mathbf{X}}_{1\text{pred}} = (\bar{\mathbf{X}}_1^{\top}, \bar{\mathbf{X}}_{1\text{new}}^{\top})^{\top}$, whereas $\bar{\Sigma}_{1\text{pred}}$ denotes the block-diagonal matrix that has $\bar{\Sigma}_{1\text{pred}\{11\}} = \bar{\Sigma}_1$ and $\bar{\Sigma}_{1\text{pred}\{22\}} = \bar{\Sigma}_{1\text{new}}$. The two quantities $\boldsymbol{\gamma}_{\text{pred}}$ and Γ_{pred} , and, implicitly, ξ_{pred} , Ω_{pred} and Δ_{pred} , are constructed analogously to the posterior parameters in Theorem 1, after replacing the original data with the enriched ones $(\bar{\mathbf{y}}_{1\text{pred}}, \bar{\mathbf{X}}_{1\text{pred}}, \bar{\Sigma}_{1\text{pred}})$ and $(\bar{\mathbf{y}}_{0\text{pred}}, \bar{\mathbf{X}}_{0\text{pred}}, \bar{\Sigma}_{0\text{pred}})$, where $\bar{\mathbf{y}}_{0\text{pred}} = (\bar{\mathbf{y}}_0^{\top}, \bar{\mathbf{y}}_{0\text{new}}^{\top})^{\top}$, $\bar{\mathbf{X}}_{0\text{pred}} = (\bar{\mathbf{X}}_0^{\top}, \bar{\mathbf{X}}_{0\text{new}}^{\top})^{\top}$ and $\bar{\Sigma}_{0\text{pred}}$ is the block-diagonal matrix with $\bar{\Sigma}_{0\text{pred}\{11\}} = \bar{\Sigma}_0$ and $\bar{\Sigma}_{0\text{pred}\{22\}} = \bar{\Sigma}_{0\text{new}}$.

Before concluding the overview of the SUN properties which facilitate posterior inference, it shall be emphasized that SUNs are also closed under marginalization, linear combinations and conditioning (Arellano-Valle and Azzalini 2006, 2021). This means, for instance, that the posterior distribution of any sub-vector $\beta_{[j]}$, with $j \subset \{1, \dots, \bar{p}\}$ is a $\text{SUN}_{|j|, \bar{n}+\bar{n}_0}(\xi_{\text{post}\{j\}}, \Omega_{\text{post}\{j\}}, \Delta_{\text{post}\{j\}}, \boldsymbol{\gamma}_{\text{post}}, \Gamma_{\text{post}})$, where $\Delta_{\text{post}\{j\}}$ corresponds to the matrix Δ_{post} after deleting all the rows with indexes not in j . Therefore, setting $j = \{j\}$ shows that the posterior of each $\beta_j, j = 1, \dots, \bar{p}$ is still a SUN. Similarly, the posterior distribution for the linear combination $(\mathbf{a} + \mathbf{A}^{\top} \beta) \in \mathbb{R}^d$ is a $\text{SUN}_{d, \bar{n}+\bar{n}_0}(\mathbf{a} + \mathbf{A}^{\top} \xi_{\text{post}}, \mathbf{A}^{\top} \Omega_{\text{post}} \mathbf{A}, [(\mathbf{A}^{\top} \Omega_{\text{post}} \mathbf{A}) \odot \mathbf{I}_d]^{-1/2} \mathbf{A}^{\top} \omega_{\text{post}} \Delta_{\text{post}}, \boldsymbol{\gamma}_{\text{post}}, \Gamma_{\text{post}})$. In particular, this implies that the posterior distribution of any linear predictor is still SUN.

The results presented in this section also clarify that, unlike for Bayesian linear regression with Gaussian priors, it is not immediate to disentangle the role of the prior parameters from the one of the data in the functionals and shape of the SUN posterior treated in this article. In fact, as clarified in (11), each of these quantities covers multiple roles in controlling location, scale and skewness; see Durante (2019) for an attempt to separate the effect of the different terms in probit regression with Gaussian priors.

4. Computational Methods

The results presented in Section 3.2 suggest that posterior inference under the models illustrated in Section 2 can be performed via closed-form solutions. This is true for any, even huge, \bar{p} as long as $\bar{n} + \bar{n}_0$ is small-to-moderate, but not when $\bar{n} + \bar{n}_0$ exceeds few hundreds (Durante 2019; Fasano and Durante 2022). In fact, (9)–(14) require evaluation of cumulative distribution functions of $(\bar{n} + \bar{n}_0)$ -variate Gaussians or sampling from $(\bar{n} + \bar{n}_0)$ -variate truncated normals, which is known to be computationally challenging in high dimensions (Genz 1992; Genz and Bretz 2009; Botev 2017; Genton, Keyes, and Turkiyyah 2018; Cao et al. 2019, 2021). This motivates still active research on developing sampling-based methods and accurate deterministic approximations for tractable Bayesian inference under the models in Section 2. Sections 4.1–4.3 review, unify, extend and compare both past and more recent developments along these lines.

4.1. Analytical Methods

As discussed above, the evaluation of high-dimensional Gaussian integrals with linear constraints, such as those found in (9)–(14), is a longstanding problem (e.g., Genz 1992; Genz and Bretz 2002; Miwa, Hayter, and Kuriki 2003; Gassmann 2003; Genz 2004; Craig 2008; Ridgway 2016; Botev 2017; Genton, Keyes, and Turkiyyah 2018; Cao et al. 2019; Gessner, Kanjilal, and Hennig 2020; Cao et al. 2021).

A popular class of strategies for evaluating these Gaussian integrals encompasses several extensions of the original separation of variables estimator initially proposed by Genz (1992). This solution recasts the problem as a sequence of tractable one-dimensional integrals, which are evaluated numerically via a randomized quasi-Monte Carlo sampling. As suggested in, for example, Genz and Bretz (2009), the variance of the resulting estimator can be further reduced by means of variable reordering. More recently, Botev (2017) proposed a new solution relying on an optimal exponential tilting of the Genz (1992) construction, which is found by solving efficiently a minimax saddle-point problem, and then used as an effective importance sampling proposal. While still providing an unbiased estimate, this technique achieves practical reduction of the estimator variance by orders of magnitude. Moreover, this procedure remains effective in settings where the Genz (1992) method cannot provide reliable estimates. Such a solution, available in the R library `TruncatedNormal`, remains generally tractable in a few hundreds of dimensions, but it progressively slows down beyond this regime. To achieve scalability in higher dimensions, recent solutions leverage low-rank hierarchical block structures of the covariance matrix within the high-dimensional Gaussian integral to decompose the problem into a sequence of smaller-dimensional ones which facilitate reduction of computational cost while preserving accuracy (Genton, Keyes, and Turkiyyah 2018; Cao et al. 2019, 2021). Among these alternatives, the one proposed in Cao et al. (2021) provides a state-of-the-art extension of the original separation of variables estimator which incorporates both an effective tile-low-rank representation of the covariance matrix and an iterative block-reordering scheme to obtain notable improvements in runtimes and scalability. For instance, such a solution has been recently adapted to the problem of evaluating predictive probabilities in high-dimensional probit Gaussian processes with \bar{n}_0 and \bar{p} in tens of thousands (Cao, Durante, and Genton 2022), obtaining remarkable improvements over state-of-the-art methods.

There are also alternative solutions beyond the classical separation of variables technique. For example, Ridgway (2016) developed a sequential Monte Carlo sampler to compute Gaussian orthant probabilities, adding a dimension at each step, combined with carefully-designed MCMC moves. More recently, Gessner, Kanjilal, and Hennig (2020) constructed an efficient estimator of Gaussian integrals with linear domain constraints, that decomposes the problem into a sequence of easier-to-solve conditional probabilities, based on nested domains. Each internal step uses an analytic version of elliptical slice sampling, exploiting the availability of closed-form solutions for the intersections between the ellipses and linear constraints. The authors reported evidence of effectiveness of such method even for thousands-dimensional integrals. Further strategies can be

found in, for example, Genz and Bretz (2002), Miwa, Hayter, and Kuriki (2003), Gassmann (2003), Genz (2004), Craig (2008), and Trinh and Genz (2015).

Interestingly, some of the aforementioned strategies also provide, as a byproduct, effective solutions for sampling from multivariate truncated normals, which can be useful to generate values from the SUN posterior via the additive representation in (11). These methods can be found, for example, in Botev (2017) and in Gessner, Kanjilal, and Hennig (2020). Motivated by inference on a phylogenetic multivariate probit model, Zhang et al. (2021) recently employed an alternative scheme for sampling from truncated normals with dimension above 10,000, via a bouncy particle sampler. See also Pakman and Paninski (2014) for an Hamiltonian Monte Carlo scheme, incorporating the truncations via hard walls and exploiting the possibility to integrate exactly the Hamiltonian equations.

All the above solutions provide effective methods for evaluating Gaussian cumulative distribution functions and, possibly, sampling from multivariate truncated normals. However, such procedures are still subject to a tradeoff between accuracy and computational tractability which is often specific to the model analyzed and to the size of the data, thereby motivating still ongoing research. Due to this, it is difficult to identify a generally-applicable gold-standard among the aforementioned techniques, although, in practice, the method by Botev (2017) has often notable performance when applied to (9)–(14) in small-to-moderate size settings with $\bar{n} + \bar{n}_0$ in the order of few hundreds. Higher-dimensional problems may require more scalable solutions (e.g., Gessner, Kanjilal, and Hennig 2020; Cao et al. 2021; Zhang et al. 2021), even if more extensive empirical analyses are required to assess these methods.

4.2. Sampling-based Methods

Whenever the interest is on more complex posterior functionals beyond those derived in Section 3.2, an effective solution is to consider Monte Carlo estimates based on samples from $p(\boldsymbol{\beta}|\mathbf{y})$. While generally-applicable MCMC strategies such as state-of-the-art implementations of Hamiltonian Monte Carlo (e.g., Hoffman and Gelman 2014) and Metropolis–Hastings (e.g., Roberts and Rosenthal 2001) can be considered, a widely-implemented class of algorithms within the context of the models presented in Section 2 are data augmentation Gibbs samplers (see, e.g., Chib 1992; Albert and Chib 1993; McCulloch and Rossi 1994; Chib and Greenberg 1998; Albert and Chib 2001; Imai and Van Dyk 2005; Holmes and Held 2006). This is because the formulations in Section 2 rely on Gaussian latent utilities which are assigned a regression model with coefficients $\boldsymbol{\beta}$. Therefore, treating these utilities as augmented data restores Gaussian-Gaussian conjugacy between the prior for $\boldsymbol{\beta}$ and the likelihood of the augmented utilities, which can be in turn sampled from independent truncated normal full-conditionals, given $\boldsymbol{\beta}$ and the censoring information provided by the observed \mathbf{y} . This yields tractable Gibbs samplers that iterate among these two steps, thus producing samples from the posterior of $\boldsymbol{\beta}$.

Although the above techniques have been proposed only for a subset of the models presented in Section 2, and in separate contributions mainly focusing on Gaussian priors (e.g.,

Chib 1992; Albert and Chib 1993; McCulloch and Rossi 1994; Holmes and Held 2006), the comprehensive framework in (1), and the general conjugacy results reported in Section 3 allow to unify these MCMC strategies within a broad construction which can be applied to any model in Section 2, even beyond those currently studied, and holds not only for Gaussian prior distributions, but also for the general SUN ones. Letting

$$\begin{aligned} \mathbf{X}_{\text{post}} &= \mathbf{\Delta}_{\text{post}}^\top \bar{\mathbf{\Omega}}_{\text{post}}^{-1} \boldsymbol{\omega}_{\text{post}}^{-1}, & \boldsymbol{\eta}_{\text{post}} &= \boldsymbol{\gamma}_{\text{post}} - \mathbf{X}_{\text{post}} \boldsymbol{\xi}_{\text{post}}, \\ \boldsymbol{\Sigma}_{\text{post}} &= \mathbf{\Gamma}_{\text{post}} - \mathbf{\Delta}_{\text{post}}^\top \bar{\mathbf{\Omega}}_{\text{post}}^{-1} \mathbf{\Delta}_{\text{post}}, \end{aligned}$$

this general Gibbs sampler can be obtained by noticing that, due to (8), the density kernel of the SUN posterior in Theorem 1 coincides with $\phi_{\bar{p}}(\boldsymbol{\beta} - \boldsymbol{\xi}_{\text{post}}; \boldsymbol{\Omega}_{\text{post}}) \Phi_{\bar{n}+\bar{n}_0}(\boldsymbol{\eta}_{\text{post}} + \mathbf{X}_{\text{post}}\boldsymbol{\beta}; \boldsymbol{\Sigma}_{\text{post}})$, where the cumulative distribution function term can be also written as $\int \phi_{\bar{n}+\bar{n}_0}(\bar{\mathbf{z}} - (\boldsymbol{\eta}_{\text{post}} + \mathbf{X}_{\text{post}}\boldsymbol{\beta}); \boldsymbol{\Sigma}_{\text{post}}) \mathbb{1}(\bar{\mathbf{z}} > \mathbf{0}) d\bar{\mathbf{z}}$. Hence, extending the augmented-data representation by Albert and Chib (1993)—see also Fasano and Durante (2022)—and leveraging standard properties of multivariate Gaussian and truncated normals, this alternative formulation implies a generally-applicable data augmentation Gibbs sampler relying on the full-conditional distributions

$$\begin{aligned} (\boldsymbol{\beta} | \bar{\mathbf{z}}) &\sim N_{\bar{p}}(\mathbf{V}_{\text{post}}[\mathbf{X}_{\text{post}}^\top \boldsymbol{\Sigma}_{\text{post}}^{-1}(\bar{\mathbf{z}} - \boldsymbol{\eta}_{\text{post}}) + \boldsymbol{\Omega}_{\text{post}}^{-1} \boldsymbol{\xi}_{\text{post}}], \mathbf{V}_{\text{post}}), \\ (\bar{\mathbf{z}} | \mathbf{y}, \boldsymbol{\beta}) &\sim \text{TN}_{\bar{n}+\bar{n}_0}(\mathbf{0}; \boldsymbol{\eta}_{\text{post}} + \mathbf{X}_{\text{post}}\boldsymbol{\beta}, \boldsymbol{\Sigma}_{\text{post}}), \end{aligned} \tag{15}$$

where $\mathbf{V}_{\text{post}} = (\boldsymbol{\Omega}_{\text{post}}^{-1} + \mathbf{X}_{\text{post}}^\top \boldsymbol{\Sigma}_{\text{post}}^{-1} \mathbf{X}_{\text{post}})^{-1}$. Hence, available Gibbs samplers for specific models within (1) and yet unexplored extensions to the whole class under general SUN priors, can be readily obtained as special cases of (15) under suitable specification of the posterior parameters defining the above full-conditionals. It shall also be emphasized that the sampling from the $(\bar{n} + \bar{n}_0)$ -dimensional truncated normal distribution in (15) is usually simplified by the conditional independence properties among the latent utilities underlying most of the regression models presented in Sections 2.1–2.4. This means that $\boldsymbol{\Sigma}_{\text{post}}$ is either diagonal or block-diagonal, often with small-dimensional blocks, and, therefore, sampling from $(\bar{\mathbf{z}} | \mathbf{y}, \boldsymbol{\beta})$ simply requires to draw values from univariate or low-dimensional truncated normals. Nonetheless, as discussed in Johndrow et al. (2018) the dependence structure between $\boldsymbol{\beta}$ and $\bar{\mathbf{z}}$ can still yield to poor mixing; see also Qin and Hobert (2019) for detailed convergence analysis.

An effective option to obviate the above mixing issues is to sample iid values from the joint posterior $p(\boldsymbol{\beta}, \bar{\mathbf{z}} | \mathbf{y})$, instead of autocorrelated ones as in (15). Extending the derivations by Holmes and Held (2006) to the whole class of models in (1), under SUN priors (8), this task can be accomplished by noting that $p(\boldsymbol{\beta}, \bar{\mathbf{z}} | \mathbf{y}) = p(\boldsymbol{\beta} | \bar{\mathbf{z}} | \mathbf{y}) p(\bar{\mathbf{z}} | \mathbf{y})$, where $p(\boldsymbol{\beta} | \bar{\mathbf{z}} | \mathbf{y})$ coincides with the density of the Gaussian in (15), whereas $p(\bar{\mathbf{z}} | \mathbf{y})$ is obtained by marginalizing out from the truncated normal in (15) the $\boldsymbol{\beta}$ vector with density $\phi_{\bar{p}}(\boldsymbol{\beta} - \boldsymbol{\xi}_{\text{post}}; \boldsymbol{\Omega}_{\text{post}})$. Leveraging standard properties of Gaussian and truncated normal random variables, and recalling Holmes and Held (2006), this implies that

$$(\bar{\mathbf{z}} | \mathbf{y}) \sim \text{TN}_{\bar{n}+\bar{n}_0}(\mathbf{0}; \boldsymbol{\gamma}_{\text{post}}, \mathbf{\Gamma}_{\text{post}}). \tag{16}$$

Replacing the full-conditional multivariate truncated normal in (15) with the one in (16), yields to a scheme for sampling iid values from $p(\boldsymbol{\beta}, \bar{\mathbf{z}} | \mathbf{y})$ and, as a direct consequence, from the

posterior $p(\boldsymbol{\beta} | \mathbf{y})$ of interest. To do this, it is sufficient to draw $\bar{\mathbf{z}}$ from (16) and then generate a value for $\boldsymbol{\beta}$ by sampling from the Gaussian in (15) with mean evaluated at the sampled value of $\bar{\mathbf{z}}$. This routine is closely related to the iid sampler based on the additive representation of the SUN in (11) that relies on a linear combination among samples from \bar{p} -variate Gaussians and $(\bar{n} + \bar{n}_0)$ -variate truncated normals (Durante 2019; Fasano et al. 2021; Fasano and Durante 2022).

Although the above strategies effectively address the mixing and convergence issues of the Gibbs sampler in (15), the multivariate truncated normal in (16) is often more challenging from a computational perspective relative to the one in (15). In fact, marginalizing out $\boldsymbol{\beta}$ in $\text{TN}_{\bar{n}+\bar{n}_0}(\mathbf{0}; \boldsymbol{\eta}_{\text{post}} + \mathbf{X}_{\text{post}}\boldsymbol{\beta}, \boldsymbol{\Sigma}_{\text{post}})$ induces dependence among the latent utilities in $\bar{\mathbf{z}}$. This means that, unlike for $\boldsymbol{\Sigma}_{\text{post}}$, the covariance matrix $\mathbf{\Gamma}_{\text{post}}$ of the truncated normal in (16) has no more a diagonal or block-diagonal structure and, hence, $p(\bar{\mathbf{z}} | \mathbf{y})$ does not factorize as the product of univariate or low-dimensional truncated normals as for $p(\bar{\mathbf{z}} | \mathbf{y}, \boldsymbol{\beta})$ in (15), making the sampling from (16) more challenging when $(\bar{n} + \bar{n}_0)$ is large. In probit regression, Holmes and Held (2006) address such issue by leveraging the closure under conditioning properties discussed, for example, in Horrace (2005) to sample iteratively from the univariate truncated normal full-conditionals $p(\bar{z}_i | \bar{\mathbf{z}}_{-i}, \mathbf{y})$, for $i = 1, \dots, \bar{n} + \bar{n}_0$. However, this strategy implies a Gibbs-sampling routine which may be still subject to mixing issues. Alternatively, it is possible to sample directly from $p(\bar{\mathbf{z}} | \mathbf{y})$ in (16) leveraging the state-of-the-art schemes presented in Section 4.1 (e.g., Botev 2017; Gessner, Kanjilal, and Hennig 2020). However, there is still the lack of a generally-applicable gold-standard for any size of \bar{p} and $\bar{n} + \bar{n}_0$.

4.3. Deterministic Approximation-based Methods

Even resorting to state-of-the-art solutions, sampling from the posterior distribution is often prohibitive for high-dimensional datasets and large sample sizes (e.g., Chopin and Ridgway 2017). In these scenarios, an effective solution is to consider deterministic approximations of the exact posterior. Sections 4.3.1–4.3.2 provide a unified treatment of classical and more recent VB (Blei, Kucukelbir, and McAuliffe 2017) and EP (Minka 2001) approximations which are widely-implemented solutions in the context of the models considered in this article; see Chopin and Ridgway (2017) for a review of alternative methods, such as Laplace approximation and INLA (Rue, Martino, and Chopin 2009). A detailed derivation and discussion of the computational costs can be found in Section S3 of the supplementary materials.

4.3.1. Variational Bayes (VB)

VB solves a constrained optimization problem that aims at finding the approximating density which is the closest, in Kullback–Leiber (KL) divergence (Kullback and Leibler 1951), to the exact posterior, among all the densities within a pre-specified tractable family facilitating Bayesian inference. Recalling Blei, Kucukelbir, and McAuliffe (2017), within the context of models admitting conditionally conjugate constructions with global parameters $\boldsymbol{\beta}$ and local augmented data $\bar{\mathbf{z}}$ —such as for the formulations in Section 2—the solution of the optimization problem often benefits from taking $p(\boldsymbol{\beta}, \bar{\mathbf{z}} | \mathbf{y})$ as the target density to be approximated,

which in turn would yield to an approximation for $p(\boldsymbol{\beta}|\mathbf{y})$ after marginalizing out $\bar{\mathbf{z}}$ (Girolami and Rogers 2006; Consonni and Marin 2007; Fasano, Durante, and Zanella 2022; Fasano and Durante 2022). As for the choice of the approximating family \mathcal{Q} , classical solutions (e.g., Girolami and Rogers 2006; Consonni and Marin 2007) rely on the mean-field assumption (e.g., Blei, Kucukelbir, and McAuliffe 2017) which can be generally expressed as $\mathcal{Q}_{\text{MF}} = \{q(\boldsymbol{\beta}, \bar{\mathbf{z}}) : q(\boldsymbol{\beta}, \bar{\mathbf{z}}) = q(\boldsymbol{\beta}) \prod_{c=1}^C q(\bar{\mathbf{z}}_c)\}$, where $\bar{\mathbf{z}}_1, \dots, \bar{\mathbf{z}}_C$ denote distinct sub-vectors of $\bar{\mathbf{z}}$, such that $\bar{\mathbf{z}} = (\bar{\mathbf{z}}_1^\top, \dots, \bar{\mathbf{z}}_C^\top)^\top$. Note that the choice of how to factorize $q(\bar{\mathbf{z}})$ in C independent blocks is often guided by the dependence structures in $\bar{\mathbf{z}}$. For instance, in models relying on conditionally independent latent utilities, such as those in Section 2, it is common to factorize $q(\bar{\mathbf{z}})$ consistent with these conditionally independent sub-vectors. In fact, as illustrated in the context of probit (e.g., Consonni and Marin 2007) and multinomial probit (e.g., Girolami and Rogers 2006), even without assuming a specific factorization for $q(\bar{\mathbf{z}})$, that is, $C = 1$, the optimum $q_{\text{MF}}^*(\bar{\mathbf{z}})$ within the class \mathcal{Q}_{MF} would still factorize as $\prod_{c=1}^C q_{\text{MF}}^*(\bar{\mathbf{z}}_c)$, where $\bar{\mathbf{z}}_1, \dots, \bar{\mathbf{z}}_C$ correspond to the subsets of conditionally independent utilities, as implied by the chosen model and prior.

Summarizing the above discussion, the mean-field variational Bayes (MF-VB) solution can be formalized as

$$\begin{aligned} q_{\text{MF}}^*(\boldsymbol{\beta}, \bar{\mathbf{z}}) &= \operatorname{argmin}_{q(\boldsymbol{\beta}, \bar{\mathbf{z}}) \in \mathcal{Q}_{\text{MF}}} \text{KL}[q(\boldsymbol{\beta}, \bar{\mathbf{z}}) \| p(\boldsymbol{\beta}, \bar{\mathbf{z}}|\mathbf{y})] \\ &= \operatorname{argmax}_{q(\boldsymbol{\beta}, \bar{\mathbf{z}}) \in \mathcal{Q}_{\text{MF}}} \text{ELBO}[q(\boldsymbol{\beta}, \bar{\mathbf{z}})], \end{aligned} \quad (17)$$

since $\text{ELBO}[q(\boldsymbol{\beta}, \bar{\mathbf{z}})] = -\text{KL}[q(\boldsymbol{\beta}, \bar{\mathbf{z}}) \| p(\boldsymbol{\beta}, \bar{\mathbf{z}}|\mathbf{y})] + \log p(\mathbf{y})$. Recalling Blei, Kucukelbir, and McAuliffe (2017), (17) can be solved by leveraging a tractable coordinate ascent variational inference (CAVI) algorithm which iteratively updates the solution of the approximating densities for both $\boldsymbol{\beta}$ and $\bar{\mathbf{z}}$ via $q_{\text{MF}}^{(t)}(\boldsymbol{\beta}) \propto \exp\{\mathbb{E}_{q(\bar{\mathbf{z}})}[\log p(\boldsymbol{\beta}|\mathbf{y}, \bar{\mathbf{z}})]\}$ and $q_{\text{MF}}^{(t)}(\bar{\mathbf{z}}_c) \propto \exp\{\mathbb{E}_{q(\boldsymbol{\beta}, \bar{\mathbf{z}}_{-c})}[\log p(\bar{\mathbf{z}}_c|\mathbf{y}, \boldsymbol{\beta}, \bar{\mathbf{z}}_{-c})]\}$, for $c = 1, \dots, C$, where $\bar{\mathbf{z}}_{-c}$ coincides with $\bar{\mathbf{z}}$ without the sub-vector $\bar{\mathbf{z}}_c$, while the expectation is taken with respect to the most recent update of the variational density over the other conditioning variables. Replacing the full-conditional distributions in these expressions with those in (15), and leveraging the closure under conditioning of multivariate truncated normals (Horrace 2005), yields a general MF-VB that extends Girolami and Rogers (2006) and Consonni and Marin (2007) to the whole class of models and priors presented in Sections 2–3, and can be obtained via closed-form CAVI updates. More specifically, let $\bar{\boldsymbol{\eta}}_{\text{post}} = \mathbb{E}_{q(\bar{\mathbf{z}})}(\bar{\boldsymbol{\eta}}) - \boldsymbol{\eta}_{\text{post}}$, where $\mathbb{E}_{q(\bar{\mathbf{z}})}(\bar{\boldsymbol{\eta}}) = [\mathbb{E}_{q(\bar{\mathbf{z}}_1)}(\bar{\boldsymbol{\eta}}_1^\top), \dots, \mathbb{E}_{q(\bar{\mathbf{z}}_C)}(\bar{\boldsymbol{\eta}}_C^\top)]^\top$, and define $\boldsymbol{\Sigma}_{\text{post}(c)} = \boldsymbol{\Sigma}_{\text{post}[c,c]} - \boldsymbol{\Sigma}_{\text{post}[c,-c]}(\boldsymbol{\Sigma}_{\text{post}[-c,-c]})^{-1}\boldsymbol{\Sigma}_{\text{post}[-c,c]}$, with $\boldsymbol{\Sigma}_{\text{post}[c,c]}$, $\boldsymbol{\Sigma}_{\text{post}[c,-c]}$, $\boldsymbol{\Sigma}_{\text{post}[-c,c]}$, and $\boldsymbol{\Sigma}_{\text{post}[-c,-c]}$, corresponding to the four blocks of $\boldsymbol{\Sigma}_{\text{post}}$ when partitioned to highlight the sub-vector $\bar{\mathbf{z}}_c$ against all the others in $\bar{\mathbf{z}}_{-c}$. Then, the CAVI updates for MF-VB are given by

$$\begin{aligned} q_{\text{MF}}^{(t)}(\boldsymbol{\beta}) &= \phi_{\bar{p}}(\boldsymbol{\beta} - \mathbf{V}_{\text{post}}(\mathbf{X}_{\text{post}}^\top \boldsymbol{\Sigma}_{\text{post}}^{-1} \bar{\boldsymbol{\eta}}_{\text{post}} + \boldsymbol{\Omega}_{\text{post}}^{-1} \boldsymbol{\xi}_{\text{post}}); \mathbf{V}_{\text{post}}), \\ q_{\text{MF}}^{(t)}(\bar{\mathbf{z}}_c) &\propto \phi_{n_c}(\bar{\mathbf{z}}_c - \mathbb{E}_{q(\boldsymbol{\beta}, \bar{\mathbf{z}}_{-c})}(\boldsymbol{\mu}_c); \boldsymbol{\Sigma}_{\text{post}(c)}) \mathbb{1}(\bar{\mathbf{z}}_c > \mathbf{0}), \\ &\quad \text{for } c = 1, \dots, C, \end{aligned} \quad (18)$$

where n_c is the dimension of $\bar{\mathbf{z}}_c$, while $\mathbb{E}_{q(\boldsymbol{\beta}, \bar{\mathbf{z}}_{-c})}(\boldsymbol{\mu}_c) = \boldsymbol{\eta}_{\text{post}[c]} + \mathbf{X}_{\text{post}[c]} \mathbb{E}_{q(\boldsymbol{\beta})}(\boldsymbol{\beta}) + \boldsymbol{\Sigma}_{\text{post}[c,-c]}(\boldsymbol{\Sigma}_{\text{post}[-c,-c]})^{-1}[\mathbb{E}_{q(\bar{\mathbf{z}}_{-c})}(\bar{\mathbf{z}}_{-c}) - \boldsymbol{\eta}_{\text{post}[-c]} - \mathbf{X}_{\text{post}[-c]} \mathbb{E}_{q(\boldsymbol{\beta})}(\boldsymbol{\beta})]$. The quantities $\mathbf{X}_{\text{post}[c]}$, $\mathbf{X}_{\text{post}[-c]}$, $\boldsymbol{\eta}_{\text{post}[c]}$ and $\boldsymbol{\eta}_{\text{post}[-c]}$ within (18) denote the rows of \mathbf{X}_{post} and

$\boldsymbol{\eta}_{\text{post}}$ corresponding to $\bar{\mathbf{z}}_c$ and $\bar{\mathbf{z}}_{-c}$, respectively. Hence, according to (18), MF-VB for the whole class of models and priors in Sections 2–3 can be implemented via a simple CAVI routine providing Gaussian and truncated normal approximating densities for $\boldsymbol{\beta}$ and $\bar{\mathbf{z}}_1, \dots, \bar{\mathbf{z}}_C$, respectively, which only require updating of the corresponding means with respect to the most recent density estimate of the other conditioning variables, until convergence of the ELBO. Computing the Gaussian expectation $\mathbb{E}_{q(\boldsymbol{\beta})}(\boldsymbol{\beta})$ poses no computational difficulties, whereas, recalling Sections 3.2 and 4.1, evaluating the mean $\mathbb{E}_{q(\bar{\mathbf{z}}_c)}(\bar{\mathbf{z}}_c)$, for $c = 1, \dots, C$ of the truncated normals may be challenging when n_c is large. Nonetheless, n_c is typically equal to 1 or to a small value when factorizing $q(\bar{\mathbf{z}})$ consistent with the diagonal block structures of $\boldsymbol{\Sigma}_{\text{post}}$ that are implied by most of the models in Sections 2.1–2.4. This means that the MF-VB solutions for the local variables $\bar{\mathbf{z}}$ correspond to tractable low-dimensional truncated normals whose expectation can be computed via efficient routines, such as the one in the R library `MomTrunc` (Galarza Morales et al. 2022).

Although MF-VB provides a scalable and widely-applicable solution under the regression models considered in this article, as shown by Fasano, Durante, and Zanella (2022) in the context of probit regression with Gaussian priors, the resulting Gaussian approximation $q_{\text{MF}}^*(\boldsymbol{\beta})$ is characterized by low accuracy, both theoretically and empirically, in high dimensions, especially when $\bar{p} > \bar{n} + \bar{n}_0$. These drawbacks are evident not only in a general underestimation of posterior uncertainty, but also in the tendency to over-shrink the locations and to induce bias in the predictive probabilities, thereby affecting the reliability of Bayesian inference under $q_{\text{MF}}^*(\boldsymbol{\beta})$. To address these fundamental issues and improve the accuracy of VB in high dimension, Fasano, Durante, and Zanella (2022) and Fasano and Durante (2022) propose a partially-factorized MF-VB solution (PFM-VB) which replaces the classical mean-field family $\mathcal{Q}_{\text{MF}} = \{q(\boldsymbol{\beta}, \bar{\mathbf{z}}) : q(\boldsymbol{\beta}, \bar{\mathbf{z}}) = q(\boldsymbol{\beta}) \prod_{c=1}^C q(\bar{\mathbf{z}}_c)\}$ with the more flexible partially-factorized one $\mathcal{Q}_{\text{PFM}} = \{q(\boldsymbol{\beta}, \bar{\mathbf{z}}) : q(\boldsymbol{\beta}, \bar{\mathbf{z}}) = q(\boldsymbol{\beta}|\bar{\mathbf{z}}) \prod_{c=1}^C q(\bar{\mathbf{z}}_c)\}$, that avoids assuming independence between $\boldsymbol{\beta}$ and $\bar{\mathbf{z}}$ as in mean-field, and only factorizes $q(\bar{\mathbf{z}})$ as $\prod_{c=1}^C q(\bar{\mathbf{z}}_c)$. The structure of this enlarged family is directly motivated by the form of the actual joint posterior $p(\boldsymbol{\beta}, \bar{\mathbf{z}}|\mathbf{y})$. In fact, as highlighted in Section 4.2, $p(\boldsymbol{\beta}, \bar{\mathbf{z}}|\mathbf{y})$ can be rewritten as $p(\boldsymbol{\beta}|\mathbf{y}, \bar{\mathbf{z}})p(\bar{\mathbf{z}}|\mathbf{y})$, where $p(\boldsymbol{\beta}|\mathbf{y}, \bar{\mathbf{z}})$ is the density of the Gaussian full-conditional in (15), whereas $p(\bar{\mathbf{z}}|\mathbf{y})$ is the one of the $(\bar{n} + \bar{n}_0)$ -variate truncated normal with full covariance matrix in (16); see also Holmes and Held (2006). Therefore, since the Gaussian form of $p(\boldsymbol{\beta}|\mathbf{y}, \bar{\mathbf{z}})$ does not seem to pose computational difficulties, it is reasonable to preserve dependence between $\boldsymbol{\beta}$ and $\bar{\mathbf{z}}$ in \mathcal{Q}_{PFM} and only approximate the intractable multivariate truncated normal density $p(\bar{\mathbf{z}}|\mathbf{y})$ via the product $\prod_{c=1}^C q(\bar{\mathbf{z}}_c)$ of low-dimensional tractable ones. In addition, when the block partitions under MF-VB and PFM-VB coincide, $\mathcal{Q}_{\text{MF}} \subset \mathcal{Q}_{\text{PFM}}$. Hence, it is guaranteed that the optimum $q_{\text{PFM}}^*(\boldsymbol{\beta}, \bar{\mathbf{z}})$ under \mathcal{Q}_{PFM} is never less accurate than $q_{\text{MF}}^*(\boldsymbol{\beta}, \bar{\mathbf{z}})$, that is, $\text{KL}[q_{\text{PFM}}^*(\boldsymbol{\beta}, \bar{\mathbf{z}}) \| p(\boldsymbol{\beta}, \bar{\mathbf{z}}|\mathbf{y})] \leq \text{KL}[q_{\text{MF}}^*(\boldsymbol{\beta}, \bar{\mathbf{z}}) \| p(\boldsymbol{\beta}, \bar{\mathbf{z}}|\mathbf{y})]$.

The improved accuracy of the PFM-VB approximation, combined with the simple solution for the optimization problem even under the new enlarged family \mathcal{Q}_{PFM} , have motivated several extensions of the original idea in Fasano, Durante, and Zanella (2022) to multinomial probit (Fasano and Durante

2022) and Gaussian processes (Cao, Durante, and Genton 2022), which can be, in fact, generalized to the whole class of models and priors presented in Sections 2–3. To clarify this result, first notice that by the chain rule of the KL divergence it follows that $\text{KL}[q(\boldsymbol{\beta}|\bar{\mathbf{z}}) \prod_{c=1}^C q(\bar{\mathbf{z}}_c) | p(\boldsymbol{\beta}, \bar{\mathbf{z}}|\mathbf{y})] = \mathbb{E}_{q(\bar{\mathbf{z}})}(\text{KL}[q(\boldsymbol{\beta}|\bar{\mathbf{z}}) | p(\boldsymbol{\beta}|\bar{\mathbf{z}}, \mathbf{y})]) + \text{KL}[\prod_{c=1}^C q(\bar{\mathbf{z}}_c) | p(\bar{\mathbf{z}}|\mathbf{y})]$, where the first nonnegative summand is equal to zero only when $q(\boldsymbol{\beta}|\bar{\mathbf{z}}) = p(\boldsymbol{\beta}|\bar{\mathbf{z}}, \mathbf{y})$. Therefore, $q_{\text{PFM}}^*(\boldsymbol{\beta}|\bar{\mathbf{z}})$ is the density of the exact Gaussian full-conditional in (15), while the minimizer of $\text{KL}[\prod_{c=1}^C q(\bar{\mathbf{z}}_c) | p(\bar{\mathbf{z}}|\mathbf{y})]$ can be readily obtained by applying the closure under conditioning properties (Horrace 2005) of the multivariate truncated normal in (16) to the CAVI equations $q_{\text{PFM}}^{(t)}(\bar{\mathbf{z}}_c) \propto \exp\{\mathbb{E}_{q(\bar{\mathbf{z}}_{-c})}[\log p(\bar{\mathbf{z}}_c|\mathbf{y}, \bar{\mathbf{z}}_{-c})]\}$, for $c = 1, \dots, C$. These results yield a scheme for obtaining $q_{\text{PFM}}^*(\boldsymbol{\beta}, \bar{\mathbf{z}})$, that is as tractable as the one of MF-VB in (18). Specifically, let $\bar{\boldsymbol{\eta}}_{\text{post}}^* = \bar{\mathbf{z}} - \boldsymbol{\eta}_{\text{post}}$ and $\boldsymbol{\Gamma}_{\text{post}(c)} = \boldsymbol{\Gamma}_{\text{post}[c,c]} - \boldsymbol{\Gamma}_{\text{post}[c,-c]}(\boldsymbol{\Gamma}_{\text{post}[-c,-c]})^{-1}\boldsymbol{\Gamma}_{\text{post}[-c,c]}$. Then, the CAVI equations for PFM-VB are

$$\begin{aligned} q_{\text{PFM}}^*(\boldsymbol{\beta}|\bar{\mathbf{z}}) &= \phi_{\bar{p}}(\boldsymbol{\beta} - \mathbf{V}_{\text{post}}(\mathbf{X}_{\text{post}}^T \boldsymbol{\Sigma}_{\text{post}}^{-1} \bar{\boldsymbol{\eta}}_{\text{post}}^* + \boldsymbol{\Omega}_{\text{post}}^{-1} \boldsymbol{\xi}_{\text{post}}); \mathbf{V}_{\text{post}}), \\ q_{\text{PFM}}^{(t)}(\bar{\mathbf{z}}_c) &\propto \phi_{n_c}(\bar{\mathbf{z}}_c - \mathbb{E}_{q(\bar{\mathbf{z}}_{-c})}(\bar{\boldsymbol{\mu}}_c); \boldsymbol{\Gamma}_{\text{post}(c)}) \mathbb{1}(\bar{\mathbf{z}}_c > \mathbf{0}), \\ &\text{for } c = 1, \dots, C, \end{aligned} \tag{19}$$

where $\mathbb{E}_{q(\bar{\mathbf{z}}_{-c})}(\bar{\boldsymbol{\mu}}_c)$ in (19) is defined as $\mathbb{E}_{q(\bar{\mathbf{z}}_{-c})}(\bar{\boldsymbol{\mu}}_c) = \boldsymbol{\gamma}_{\text{post}[c]} + \boldsymbol{\Gamma}_{\text{post}[c,-c]}(\boldsymbol{\Gamma}_{\text{post}[-c,-c]})^{-1}(\mathbb{E}_{q(\bar{\mathbf{z}}_{-c})}(\bar{\mathbf{z}}_{-c}) - \boldsymbol{\gamma}_{\text{post}[-c]})$, with the expectation taken with respect to the most recent density estimate of the conditioning variables, whereas the indexing of subvectors and matrix blocks is the same as the one detailed in (18).

As for the MF-VB scheme in (18), also the CAVI for PFM-VB simply requires to update the mean vectors until convergence of the ELBO. However, unlike for (18), such a scheme is only required for the truncated normal components, whereas the solution $q_{\text{PFM}}^*(\boldsymbol{\beta}|\bar{\mathbf{z}})$ is already known to coincide with $p(\boldsymbol{\beta}|\bar{\mathbf{z}}, \mathbf{y})$. This gain comes at the cost that, unlike for MF-VB, the approximation $q_{\text{PFM}}^*(\boldsymbol{\beta})$ of interest is not available as a direct output of (19). Recalling, Fasano, Durante, and Zanella (2022) and Fasano and Durante (2022), this apparent drawback can be easily addressed after noticing that, by (19) and $\bar{\boldsymbol{\eta}}_{\text{post}}^* = \bar{\mathbf{z}} - \boldsymbol{\eta}_{\text{post}}$, $q_{\text{PFM}}^*(\boldsymbol{\beta})$ is the density of the random variable distributed as a linear combination between a Gaussian, with mean $\mathbf{V}_{\text{post}}(-\mathbf{X}_{\text{post}}^T \boldsymbol{\Sigma}_{\text{post}}^{-1} \boldsymbol{\eta}_{\text{post}} + \boldsymbol{\Omega}_{\text{post}}^{-1} \boldsymbol{\xi}_{\text{post}})$ and covariance matrix \mathbf{V}_{post} , and a random vector $\bar{\mathbf{z}}$ whose joint density is approximated via the product of low-dimensional truncated normals under the CAVI updates in (19). Recalling (11), this construction coincides with the additive representation of a $\text{SUN}_{\bar{p}, \bar{n} + \bar{n}_0}$ variable that, unlike for the exact SUN posterior in Theorem 1, relies on a block-diagonal matrix $\boldsymbol{\Gamma}_{\text{PFM}}$ with C low-dimensional $n_c \times n_c$ blocks, for $c = 1, \dots, C$. This means that the computational challenges for closed-form inference under the exact SUN posterior discussed in Sections 3.2 and 4.1 are no more present for the optimal SUN approximating density $q_{\text{PFM}}^*(\boldsymbol{\beta})$, since the $(\bar{n} + \bar{n}_0)$ -variate Gaussian cumulative distribution functions and truncated normals in (9)–(14) now factorize as C low-dimensional components that can be effectively evaluated whenever n_1, \dots, n_C are small-to-moderate. For example,

$$\begin{aligned} \mathbb{E}_{\text{PFM}}(\boldsymbol{\beta}) &= \mathbf{V}_{\text{post}}(\mathbf{X}_{\text{post}}^T \boldsymbol{\Sigma}_{\text{post}}^{-1} (\mathbb{E}_{\text{PFM}}(\bar{\mathbf{z}}) - \boldsymbol{\eta}_{\text{post}}) + \boldsymbol{\Omega}_{\text{post}}^{-1} \boldsymbol{\xi}_{\text{post}}), \\ \text{var}_{\text{PFM}}(\boldsymbol{\beta}) &= \mathbf{V}_{\text{post}} + \mathbf{V}_{\text{post}} \mathbf{X}_{\text{post}}^T \boldsymbol{\Sigma}_{\text{post}}^{-1} \text{var}_{\text{PFM}}(\bar{\mathbf{z}}) \boldsymbol{\Sigma}_{\text{post}}^{-1} \mathbf{X}_{\text{post}} \mathbf{V}_{\text{post}}, \end{aligned} \tag{20}$$

where $\mathbb{E}_{\text{PFM}}(\bar{\mathbf{z}}) = [\mathbb{E}_{\text{PFM}}(\bar{\mathbf{z}}_1^T), \dots, \mathbb{E}_{\text{PFM}}(\bar{\mathbf{z}}_C^T)]^T$ comprises the expectation of each low-dimensional sub-vector $\bar{\mathbf{z}}_c$, $c = 1, \dots, C$ with respect to its optimal truncated normal approximating density, while $\text{var}_{\text{PFM}}(\bar{\mathbf{z}})$ is a block-diagonal covariance matrix with generic block $\text{var}_{\text{PFM}}(\bar{\mathbf{z}})_{[c,c]}$ denoting the covariance matrix of $\bar{\mathbf{z}}_c$ according to its optimal truncated normal approximation. As previously mentioned, each of these quantities can be effectively evaluated in small-to-moderate dimensions via, for example, the R library `MomTrunc` (Galarza Morales et al. 2022). Recalling Fasano, Durante, and Zanella (2022), the computational complexity of PFM-VB is the same as the one for MF-VB, although the new partially-factorized solution yields improved accuracy both in theory and in practice. For instance, the authors prove that, unlike for MF-VB, the KL divergence between the PFM-VB approximation and the exact posterior goes to 0 as $\bar{p} \rightarrow \infty$ for any fixed sample size, thereby providing accurate inference in high-dimensional settings at a much lower computational cost than the exact solution.

4.3.2. Expectation-Propagation (EP)

EP (Minka 2001) provides another well-established procedure for constructing a global approximation $q_{\text{EP}}^*(\boldsymbol{\beta})$ of the posterior distribution $p(\boldsymbol{\beta}|\mathbf{y})$ (see, e.g., Chopin and Ridgway 2017; Riihimäki, Jylänki, and Vehtari 2014; Vehtari et al. 2020), which often yields improved accuracy in practice, relative to VB. Contrarily to the mean-field VB methods presented in Section 4.3.1—which only impose factorized structures for the approximating densities without necessarily assuming a functional form—EP postulates that the target posterior density itself can be written as a product of factors, also referred to as sites, and then iteratively approximates each one with an element of a given family of distributions, typically Gaussian for continuous variables or multinomial for discrete ones. Moreover, in the EP scheme each update is driven by the minimization of a suitable reverse KL, instead of the forward KL as in VB. This operation tends to improve accuracy (e.g., Chopin and Ridgway 2017) and becomes particularly convenient when the approximating density $q_{\text{EP}}(\boldsymbol{\beta})$ belongs to the exponential family, since it simply requires suitable moment matching strategies between $q_{\text{EP}}(\boldsymbol{\beta})$ and $p(\boldsymbol{\beta}|\mathbf{y})$ (see, e.g., Vehtari et al. 2020; Bishop 2006, chap. 10).

Current implementations of EP for probit (Chopin and Ridgway 2017) and multinomial probit (Riihimäki, Jylänki, and Vehtari 2014) suggest that these strategies may yield practical gains for the whole class of models in Section 2, thus, motivating the development of a broadly-applicable unified EP scheme, that is unavailable to date. This section aims at covering such a gap, while providing novel closed-form expressions for moment matching of Gaussian sites leveraging the SUN conjugacy in Section 3, which also yields additional supporting arguments on the accuracy of EP for the models in Section 2.

To address such a goal, first notice that, although the likelihood in (1) is general, all the relevant examples discussed in Section 2 admit a factorized form $\prod_{c=1}^C \Phi_{\bar{n}_c}(\bar{\mathbf{y}}_{0[c]} + \bar{\mathbf{X}}_{0[c]}\boldsymbol{\beta}; \bar{\boldsymbol{\Sigma}}_{0[c,c]})$ for the intractable quantity $\Phi_{\bar{n}_0}(\bar{\mathbf{y}}_0 + \bar{\mathbf{X}}_0\boldsymbol{\beta}; \bar{\boldsymbol{\Sigma}}_0)$, where $\bar{\mathbf{y}}_0 = (\bar{\mathbf{y}}_{0[1]}^T, \dots, \bar{\mathbf{y}}_{0[C]}^T)^T$, $\bar{\mathbf{X}}_0 = (\bar{\mathbf{X}}_{0[1]}^T, \dots, \bar{\mathbf{X}}_{0[C]}^T)^T$ and $\bar{\boldsymbol{\Sigma}}_0$ is a block-diagonal matrix with generic block $\bar{\boldsymbol{\Sigma}}_{0[c,c]}$, for every $c = 1, \dots, C$. As discussed in Sections 4.2 and 4.3.1, this factorization is implied by the conditional independence among the latent

utilities, which yields to tractable one-dimensional (e.g., probit and tobit) or low-dimensional (e.g., multinomial probit) factors $\Phi_{\bar{n}_c}(\bar{\mathbf{y}}_{0[c]} + \bar{\mathbf{X}}_{0[c]}\boldsymbol{\beta}; \bar{\boldsymbol{\Sigma}}_{0[c,c]})$. Therefore, under these models, the likelihood in (1) is equal to

$$p(\mathbf{y}|\boldsymbol{\beta}) = p(\bar{\mathbf{y}}_1|\boldsymbol{\beta})p(\bar{\mathbf{y}}_0|\boldsymbol{\beta}) \tag{21}$$

$$\propto \phi_{\bar{n}_1}(\bar{\mathbf{y}}_1 - \bar{\mathbf{X}}_1\boldsymbol{\beta}; \bar{\boldsymbol{\Sigma}}_1) \prod_{c=1}^C \Phi_{\bar{n}_c}(\bar{\mathbf{y}}_{0[c]} + \bar{\mathbf{X}}_{0[c]}\boldsymbol{\beta}; \bar{\boldsymbol{\Sigma}}_{0[c,c]}),$$

thus, providing a general factorized structure that motivates EP. For ease of notation and presentation, this routine is first derived below under the Gaussian prior $p(\boldsymbol{\beta}) = \phi_{\bar{p}}(\boldsymbol{\beta} - \boldsymbol{\xi}; \boldsymbol{\Omega})$, and subsequently extended to the general class of SUN prior distributions. Updating $p(\boldsymbol{\beta}) = \phi_{\bar{p}}(\boldsymbol{\beta} - \boldsymbol{\xi}; \boldsymbol{\Omega})$ with the likelihood in (21) yields the posterior $p(\boldsymbol{\beta}|\mathbf{y})$ which can be more conveniently re-expressed as

$$p(\boldsymbol{\beta}|\mathbf{y}) \propto p(\boldsymbol{\beta}|\bar{\mathbf{y}}_1)p(\bar{\mathbf{y}}_0|\boldsymbol{\beta})$$

$$\propto \phi_{\bar{p}}(\boldsymbol{\beta} - \boldsymbol{\xi}_{\text{post}}; \boldsymbol{\Omega}_{\text{post}}) \prod_{c=1}^C \Phi_{\bar{n}_c}(\bar{\mathbf{y}}_{0[c]} + \bar{\mathbf{X}}_{0[c]}\boldsymbol{\beta}; \bar{\boldsymbol{\Sigma}}_{0[c,c]})$$

$$= l_0(\boldsymbol{\beta}) \prod_{c=1}^C l_c(\boldsymbol{\beta}) = \prod_{c=0}^C l_c(\boldsymbol{\beta}), \tag{22}$$

where $l_c(\boldsymbol{\beta}) = \Phi_{\bar{n}_c}(\bar{\mathbf{y}}_{0[c]} + \bar{\mathbf{X}}_{0[c]}\boldsymbol{\beta}; \bar{\boldsymbol{\Sigma}}_{0[c,c]})$, $c = 1, \dots, C$, correspond to the Gaussian cumulative distribution function terms in likelihood (21), whereas $l_0(\boldsymbol{\beta}) = p(\boldsymbol{\beta}|\bar{\mathbf{y}}_1) = \phi_{\bar{p}}(\boldsymbol{\beta} - \boldsymbol{\xi}_{\text{post}}; \boldsymbol{\Omega}_{\text{post}})$ is the conditional density obtained by updating the Gaussian prior $\phi_{\bar{p}}(\boldsymbol{\beta} - \boldsymbol{\xi}; \boldsymbol{\Omega})$ for $\boldsymbol{\beta}$ with the tractable factor $\phi_{\bar{n}_1}(\bar{\mathbf{y}}_1 - \bar{\mathbf{X}}_1\boldsymbol{\beta}; \bar{\boldsymbol{\Sigma}}_1)$ in likelihood (21). As a direct consequence of the results in Section 3.2, this conditional density can be obtained in closed form and coincides with the one of a Gaussian $N_{\bar{p}}(\boldsymbol{\xi}_{\text{post}}; \boldsymbol{\Omega}_{\text{post}})$ having parameters defined as in Theorem 1. Such a density acts as an intermediate prior in (22) to be updated with the intractable likelihood terms for obtaining the posterior $p(\boldsymbol{\beta}|\mathbf{y})$.

Recalling, for instance, Vehtari et al. (2020), EP approximates the above posterior with a density $q_{\text{EP}}(\boldsymbol{\beta})$ that has the same factorized form of $p(\boldsymbol{\beta}|\mathbf{y})$ in (22), and is made of $C + 1$ Gaussian sites. Hence

$$q_{\text{EP}}(\boldsymbol{\beta}) \propto \prod_{c=0}^C q_c(\boldsymbol{\beta}) = \prod_{c=0}^C \exp(-0.5\boldsymbol{\beta}^T \mathbf{Q}_c \boldsymbol{\beta} + \boldsymbol{\beta}^T \mathbf{r}_c)$$

$$= \exp(-0.5\boldsymbol{\beta}^T \mathbf{Q}_{\text{EP}} \boldsymbol{\beta} + \boldsymbol{\beta}^T \mathbf{r}_{\text{EP}}), \tag{23}$$

where \mathbf{r}_c and \mathbf{Q}_c define the natural parameters associated with the local Gaussian site c , for each $c = 0, \dots, C$, while $\mathbf{r}_{\text{EP}} = \sum_{c=0}^C \mathbf{r}_c$ and $\mathbf{Q}_{\text{EP}} = \sum_{c=0}^C \mathbf{Q}_c$ denote those of the Gaussian EP approximation $q_{\text{EP}}(\boldsymbol{\beta})$ for $p(\boldsymbol{\beta}|\mathbf{y})$. Consistent with the above expressions, the ideal goal of EP would be to obtain the optimal \mathbf{r}_{EP}^* and \mathbf{Q}_{EP}^* such that the induced Gaussian density $q_{\text{EP}}^*(\boldsymbol{\beta})$ under (23) is as close as possible to the exact $p(\boldsymbol{\beta}|\mathbf{y})$ in (22) under the reverse KL divergence $\text{KL}[p(\boldsymbol{\beta}|\mathbf{y})||q_{\text{EP}}(\boldsymbol{\beta})]$. Recalling Bishop (2006, chap. 10), the solution of this optimization problem relies on a simple moment matching, which implies that $\mathbf{r}_{\text{EP}}^* = [\text{var}(\boldsymbol{\beta}|\mathbf{y})]^{-1} \mathbb{E}(\boldsymbol{\beta}|\mathbf{y})$ and $\mathbf{Q}_{\text{EP}}^* = [\text{var}(\boldsymbol{\beta}|\mathbf{y})]^{-1}$, or, alternatively, $\boldsymbol{\xi}_{\text{EP}}^* = \mathbb{E}(\boldsymbol{\beta}|\mathbf{y})$ and $\boldsymbol{\Omega}_{\text{EP}}^* = \text{var}(\boldsymbol{\beta}|\mathbf{y})$, where $\boldsymbol{\xi}_{\text{EP}}^*$ and $\boldsymbol{\Omega}_{\text{EP}}^*$ denote the mean vector and the covariance matrix of the Gaussian EP approximation. As discussed in Section 3.2, the exact posterior is a SUN, and, hence, computing the associated moments is computationally challenging in general settings. In fact, such computational bottlenecks are those motivating the approximate schemes in Section 4.3.

To circumvent the aforementioned issue, EP relies on an iterative scheme which progressively improves $\mathbf{r}_{\text{EP}} = \sum_{c=0}^C \mathbf{r}_c$ and

$\mathbf{Q}_{\text{EP}} = \sum_{c=0}^C \mathbf{Q}_c$ by sequentially updating each term $(\mathbf{r}_c, \mathbf{Q}_c)$, for $c = 1, \dots, C$, keeping fixed the others at their previous estimate (see, e.g., Vehtari et al. 2020). Let $l_{(-c)}(\boldsymbol{\beta}) = \prod_{c' \neq c} l_{c'}(\boldsymbol{\beta})$ and $q_{(-c)}(\boldsymbol{\beta}) = \prod_{c' \neq c} q_{c'}(\boldsymbol{\beta})$ denote the product among the factors in (22)–(23), respectively, excluding the c th one. Then, EP proceeds by optimizing, for every site c , a more tractable approximation for the reverse KL divergence in which the exact posterior $p(\boldsymbol{\beta}|\mathbf{y}) \propto l_{(-c)}(\boldsymbol{\beta})l_c(\boldsymbol{\beta})$ in (22) is replaced by the intermediate hybrid density defined as $p^{(t_c)}(\boldsymbol{\beta}|\mathbf{y}) \propto q_{(-c)}^{(t_c)}(\boldsymbol{\beta})l_c(\boldsymbol{\beta})$, where t_c is the step of the algorithm which updates the site c at the t th iteration. Employing this alternative density yields a more tractable updating scheme since, by (23), $q_{(-c)}^{(t_c)}(\boldsymbol{\beta})$ is the kernel of a multivariate Gaussian distribution with natural parameters $\mathbf{r}_{\text{EP}(-c)}^{(t_c)}$ and $\mathbf{Q}_{\text{EP}(-c)}^{(t_c)}$ corresponding to $\mathbf{r}_{\text{EP}} - \mathbf{r}_c$ and $\mathbf{Q}_{\text{EP}} - \mathbf{Q}_c$, respectively, when \mathbf{r}_{EP} , \mathbf{r}_c , \mathbf{Q}_{EP} , and \mathbf{Q}_c are fixed at their most recent estimate. Therefore, $p^{(t_c)}(\boldsymbol{\beta}|\mathbf{y})$ has a single Gaussian cumulative distribution function term $l_c(\boldsymbol{\beta}) = \Phi_{\bar{n}_c}(\bar{\mathbf{y}}_{0[c]} + \bar{\mathbf{X}}_{0[c]}\boldsymbol{\beta}; \bar{\boldsymbol{\Sigma}}_{0[c,c]})$. Adapting the results in Section 3, this yields the hybrid density $p^{(t_c)}(\boldsymbol{\beta}|\mathbf{y}) \propto q_{(-c)}^{(t_c)}(\boldsymbol{\beta})l_c(\boldsymbol{\beta})$ with

$$q_{(-c)}^{(t_c)}(\boldsymbol{\beta}) = \phi_{\bar{p}}(\boldsymbol{\beta} - (\mathbf{Q}_{\text{EP}(-c)}^{(t_c)})^{-1} \mathbf{r}_{\text{EP}(-c)}^{(t_c)}; (\mathbf{Q}_{\text{EP}(-c)}^{(t_c)})^{-1}),$$

$$l_c(\boldsymbol{\beta}) = \Phi_{\bar{n}_c}(\bar{\mathbf{y}}_{0[c]} + \bar{\mathbf{X}}_{0[c]}\boldsymbol{\beta}; \bar{\boldsymbol{\Sigma}}_{0[c,c]}),$$

which implies that $p^{(t_c)}(\boldsymbol{\beta}|\mathbf{y})$ coincides with the density of the $\text{SUN}_{\bar{p}, \bar{n}_c}(\boldsymbol{\xi}_c, \boldsymbol{\Omega}_c, \boldsymbol{\Delta}_c, \boldsymbol{\gamma}_c, \boldsymbol{\Gamma}_c)$ with

$$\boldsymbol{\xi}_c = (\mathbf{Q}_{\text{EP}(-c)}^{(t_c)})^{-1} \mathbf{r}_{\text{EP}(-c)}^{(t_c)}, \quad \boldsymbol{\Omega}_c = (\mathbf{Q}_{\text{EP}(-c)}^{(t_c)})^{-1},$$

$$\boldsymbol{\Delta}_c = \bar{\boldsymbol{\Omega}}_c \boldsymbol{\omega}_c \bar{\mathbf{X}}_{0[c]}^T \mathbf{s}_c^{-1}, \quad \boldsymbol{\gamma}_c = \mathbf{s}_c^{-1} (\bar{\mathbf{y}}_{0[c]} + \bar{\mathbf{X}}_{0[c]}\boldsymbol{\xi}_c),$$

$$\boldsymbol{\Gamma}_c = \mathbf{s}_c^{-1} (\bar{\boldsymbol{\Sigma}}_{0[c,c]} + \bar{\mathbf{X}}_{0[c]}\boldsymbol{\Omega}_c \bar{\mathbf{X}}_{0[c]}^T) \mathbf{s}_c^{-1},$$

where $\mathbf{s}_c = [(\bar{\boldsymbol{\Sigma}}_{0[c,c]} + \bar{\mathbf{X}}_{0[c]}\boldsymbol{\Omega}_c \bar{\mathbf{X}}_{0[c]}^T) \odot \mathbf{I}_{\bar{n}_c}]^{1/2}$. Therefore, unlike for the exact SUN posterior, this hybrid SUN is much more tractable since the dimension of the cumulative distribution function term is \bar{n}_c , and not $\sum_{c=1}^C \bar{n}_c$ as in $p(\boldsymbol{\beta}|\mathbf{y})$. In fact, as previously discussed, \bar{n}_c is either equal to 1 or to a low value under most of the models outlined in Sections 2.1–2.4. This means that inference under the SUN with density $p^{(t_c)}(\boldsymbol{\beta}|\mathbf{y})$ can be performed via the closed-form expressions in Section 3.2, which can be effectively evaluated when \bar{n}_c is small; see also Section 4.1. In particular, it is possible to compute the expectation $\mathbb{E}^{(t_c)}(\boldsymbol{\beta}|\mathbf{y})$ and variance $\text{var}^{(t_c)}(\boldsymbol{\beta}|\mathbf{y})$ of $\boldsymbol{\beta}$ with respect to the hybrid density $p^{(t_c)}(\boldsymbol{\beta}|\mathbf{y})$ via expressions (10) evaluated at the parameters $\boldsymbol{\xi}_c, \boldsymbol{\Omega}_c, \boldsymbol{\Delta}_c, \boldsymbol{\gamma}_c$, and $\boldsymbol{\Gamma}_c$. Alternatively, leveraging the additive representation of the SUN in (11), it follows that

$$\mathbb{E}^{(t_c)}(\boldsymbol{\beta}|\mathbf{y}) = \boldsymbol{\xi}_c + \boldsymbol{\omega}_c \boldsymbol{\Delta}_c \boldsymbol{\Gamma}_c^{-1} \mathbb{E}(\mathbf{U}_{1c}) = \boldsymbol{\xi}_{\text{EP}}^{(t_c)},$$

$$\text{var}^{(t_c)}(\boldsymbol{\beta}|\mathbf{y}) = \boldsymbol{\Omega}_c - \boldsymbol{\omega}_c \boldsymbol{\Delta}_c \boldsymbol{\Gamma}_c^{-1} \boldsymbol{\Delta}_c^T \boldsymbol{\omega}_c$$

$$+ \boldsymbol{\omega}_c \boldsymbol{\Delta}_c \boldsymbol{\Gamma}_c^{-1} \text{var}(\mathbf{U}_{1c}) \boldsymbol{\Gamma}_c^{-1} \boldsymbol{\Delta}_c^T \boldsymbol{\omega}_c = \boldsymbol{\Omega}_{\text{EP}}^{(t_c)}, \tag{24}$$

where $\mathbf{U}_{1c} \sim \text{TN}_{\bar{n}_c}(-\boldsymbol{\gamma}_c; \mathbf{0}, \boldsymbol{\Gamma}_c)$ is a low-dimensional truncated normal whose expectation $\mathbb{E}(\mathbf{U}_{1c})$ and variance $\text{var}(\mathbf{U}_{1c})$ can be effectively computed via R library `MomTrunc` (Galarza Morales et al. 2022), due to the small value of \bar{n}_c . This implies that the reverse KL can be easily optimized via moment matching when $p(\boldsymbol{\beta}|\mathbf{y})$ is replaced by $p^{(t_c)}(\boldsymbol{\beta}|\mathbf{y})$, thereby obtaining the updated estimates $\mathbf{r}_{\text{EP}}^{(t_c)}$ and $\mathbf{Q}_{\text{EP}}^{(t_c)}$ for the parameters of interest \mathbf{r}_{EP} and \mathbf{Q}_{EP} at step t_c , defined as

$$\mathbf{r}_{\text{EP}}^{(t_c)} = (\boldsymbol{\Omega}_{\text{EP}}^{(t_c)})^{-1} \boldsymbol{\xi}_{\text{EP}}^{(t_c)}, \quad \mathbf{Q}_{\text{EP}}^{(t_c)} = (\boldsymbol{\Omega}_{\text{EP}}^{(t_c)})^{-1}.$$

Concurrently, the updated parameters at the site c —which are required for the subsequent steps—are $\mathbf{r}_c^{(t_c)} = \mathbf{r}_{\text{EP}}^{(t_c)} - \mathbf{r}_{\text{EP}(-c)}^{(t_c)}$ and $\mathbf{Q}_c^{(t_c)} = \mathbf{Q}_{\text{EP}}^{(t_c)} - \mathbf{Q}_{\text{EP}(-c)}^{(t_c)}$.

The above scheme is iterated multiple times $t \in \{1, 2, \dots\}$ and for each site $c = 1, \dots, C$, until convergence to a stationary point. Note that in this routine site $c = 0$ does not require to be updated sequentially. Recalling Chopin and Ridgway (2017) and Vehtari et al. (2020), $l_0(\boldsymbol{\beta})$ corresponds to the tractable Gaussian density in (22) and, hence, this term can be analytically matched to $q_0(\boldsymbol{\beta})$ in (23), obtaining $\mathbf{r}_0 = \boldsymbol{\Omega}_{\text{post}}^{-1} \boldsymbol{\xi}_{\text{post}}$ and $\mathbf{Q}_0 = \boldsymbol{\Omega}_{\text{post}}^{-1}$, where $\boldsymbol{\xi}_{\text{post}}$ and $\boldsymbol{\Omega}_{\text{post}}$ are defined as in Theorem 1. We shall also emphasize that the aforementioned EP scheme can yield, as a direct by-product, an approximation of the marginal likelihood $p(\mathbf{y})$. A detailed presentation of the step-by-step procedure to obtain such an estimate can be found in Appendix E of Vehtari et al. (2020), which shows that a key condition to compute such an approximation is the availability of the normalizing constant for the hybrid density $p^{(t_c)}(\boldsymbol{\beta}|\mathbf{y})$. Interestingly, this quantity is available in closed form for the EP scheme discussed above since $p^{(t_c)}(\boldsymbol{\beta}|\mathbf{y})$ is the density of a SUN $_{\bar{p}, \bar{n}_c}(\boldsymbol{\xi}_c, \boldsymbol{\Omega}_c, \boldsymbol{\Delta}_c, \boldsymbol{\gamma}_c, \boldsymbol{\Gamma}_c)$. Hence, recalling Section 3, its normalizing constant is $\Phi_{\bar{n}_c}(\boldsymbol{\gamma}_c; \boldsymbol{\Gamma}_c)$. Since \bar{n}_c is small, also this quantity can be effectively evaluated using, for example, the R library `TruncatedNormal` (Botev 2017).

Although EP is often more accurate than VB, it shall be noted that state-of-the-art implementations build on weaker theoretical guarantees (e.g., Bishop 2006; Chopin and Ridgway 2017; Vehtari et al. 2020) and, as discussed in the supplementary materials, are more computationally demanding. In particular, even the new EP implementation proposed in Section S3—which noticeably reduces the currently reported per-iteration cost in probit regression from $\mathcal{O}(np^2) = \mathcal{O}(\bar{n}_0 \bar{p}^2)$ (Chopin and Ridgway 2017) to $\mathcal{O}(np \cdot \min\{n, p\}) = \mathcal{O}(\bar{n}_0 \bar{p} \cdot \min\{\bar{n}_0, \bar{p}\})$ —is still not competitive with the $\mathcal{O}(n \cdot \min\{n, p\}) = \mathcal{O}(\bar{n}_0 \cdot \min\{\bar{n}_0, \bar{p}\})$ cost of MF-VB and PFM-VB (Fasano, Durante, and Zanella 2022). Moreover, there is no guarantee that the EP solution minimizes the global reverse KL $[p(\boldsymbol{\beta}|\mathbf{y})||q_{\text{EP}}(\boldsymbol{\beta})]$, nor that the routine always converges in general. Nonetheless, empirical evidence typically reports remarkable EP accuracy, which is also confirmed by the simulations in Section 5. Recalling Bishop (2006) an intuition for this notable performance is that, at each EP iteration, the sites are updated to be most accurate in regions of high posterior probability. Dehaene and Barthelmé (2018) provide more formal arguments, which show that in asymptotic settings the discrepancy between the EP solution and the exact posterior goes to 0 faster than, for instance, Laplace approximation. These results are intimately related to the log-concavity of the target posterior. Interestingly, as shown in Arellano-Valle and Azzalini (2021, sec. 3.1), SUNs are log-concave. Hence, the results in Section 3 also provide further support to EP under the models in Section 2.

Before concluding the analysis of EP, notice that extending the above derivations to the case of a more general SUN prior poses no conceptual difficulties. In fact, the Gaussian density and distribution functions appearing in the prior in (8) can be disentangled and treated as two separate sites appearing in the factorized target. In particular, the first exact site $l_0(\boldsymbol{\beta}) = \phi_{\bar{p}}(\boldsymbol{\beta} - \boldsymbol{\xi}_{\text{post}}; \boldsymbol{\Omega}_{\text{post}})$ remains unchanged, as it still arises from the combination of the Gaussian density $\phi_{\bar{p}}(\boldsymbol{\beta} - \boldsymbol{\xi}; \boldsymbol{\Omega})$ in the

prior and the likelihood term $\phi_{\bar{n}_1}(\bar{\mathbf{y}}_1 - \bar{\mathbf{X}}_1 \boldsymbol{\beta}; \bar{\boldsymbol{\Sigma}}_1)$. Conversely, the distribution function term in the prior can be simply addressed by adding a site $l_{C+1}(\boldsymbol{\beta}) = \Phi_{\bar{n}}(\boldsymbol{\gamma} + \boldsymbol{\Delta}^T \bar{\boldsymbol{\Omega}}^{-1} \boldsymbol{\omega}^{-1}(\boldsymbol{\beta} - \boldsymbol{\xi}); \boldsymbol{\Gamma} - \boldsymbol{\Delta}^T \bar{\boldsymbol{\Omega}}^{-1} \boldsymbol{\Delta})$, to be approximated via an extra term $q_{C+1}(\boldsymbol{\beta})$ in (23). As such, the only hindrances might arise from the computation of the moments for the hybrid distribution with kernel $l_{C+1}(\boldsymbol{\beta}) \prod_{c=0}^C q_c^{(t_{C+1})}(\boldsymbol{\beta})$, which still corresponds to a SUN of dimensions \bar{p} and \bar{n} . However, as mentioned in Section 3.1, the SUN prior often relies on a low \bar{n} , or alternatively factorizes as the product of independent skew-normals. In the former case, computations in low dimensions remain feasible, while in the latter case the $(C + 1)$ th exact site can be further disentangled as the product of \bar{p} sites, each one involving a tractable univariate Gaussian cumulative distribution function. This flexibility in choosing the factorization of the target distribution is a general characteristic of EP. The two extreme cases correspond, respectively, to considering the target as one single site or decomposing it into the finest factorization allowed by its analytical formulation. Any intermediate situation leads to a valid EP routine, as described above, possibly with a tradeoff between accuracy of the resulting approximation and complexity of the required computations. Finer factorizations lead to simpler moment matching calculations. Conversely, coarser factorizations yield more accurate, but expensive, approximations.

5. Empirical Studies

Insightful empirical assessments of the methods in Sections 3–4, under selected regression models, can be found in Chopin and Ridgway (2017), Durante (2019), Fasano and Durante (2022), Cao, Durante, and Genton (2022), Fasano, Durante, and Zanella (2022), Fasano et al. (2021), and Benavoli, Azzimonti, and Piga (2020, 2021); refer also to the GitHub repositories `ProbitSUN`, `PredProbitGP`, `Probit-PFMVB`, and `Dynamic-Probit-PFMVB`. These studies encompass analyses of probit regression, multinomial probit, dynamic probit, probit Gaussian processes, skewed Gaussian processes and possible combinations of these constructions, but do not cover tobit regression for which SUN conjugacy has been established in the present article and, hence, the practical consequences of this result and the associated computational methods remain unexplored.

To address such a gap, we provide empirical evidence for the performance of the computational methods in Section 4, focusing on standard tobit regression as in (7). In accomplishing this goal, we simulate a total of $n = n_0 + n_1 = 200$ observations from a tobit model, under three different proportions of censored observations $\kappa = n_0/n \in \{0.15, 0.50, 0.85\}$. This choice allows to cover a broad spectrum of scenarios which ranges from a model more similar to a Gaussian linear regression, when $\kappa = 0.15$, to one closely mimicking an unbalanced probit regression, when $\kappa = 0.85$. The p unit-specific predictors in \mathbf{x}_i , $i = 1, \dots, n$, are instead simulated from standard Gaussians, except for the intercept term, whereas the regression coefficients in $\boldsymbol{\beta}$ are generated from a uniform distribution in the range $[-5, 5]$. Exploiting the latent utility interpretation of tobit regression, the responses y_i , $i = 1, \dots, n$ are obtained by first simulating the associated utilities z_i , $i = 1, \dots, n$

Table 1. Runtimes, in seconds, of two strategies to sample 5000 realizations from the posterior distribution in tobit regression with $n = 200$.

Censoring	Method	p							
		10	20	50	100	200	400	800	1200
$\kappa = 0.85$	NUTS	2.32	4.61	9.87	26.17	50.21	134.95	323.73	659.25
	iid	2.94	2.85	18.97	15.98	7.71	3.55	4.94	9.26
$\kappa = 0.50$	NUTS	1.62	2.04	3.78	10.90	39.06	110.73	546.07	2128.08
	iid	1.30	1.21	1.20	1.84	2.58	2.33	4.21	8.25
$\kappa = 0.15$	NUTS	0.91	1.18	2.27	4.60	20.08	87.10	619.58	1973.92
	iid	0.13	0.10	0.14	0.23	0.56	1.13	2.91	6.91

NOTE: (NUTS): `rstan` implementation of No-U-Turn HMC sampler. (iid): iid sampling from the exact SUN posterior via (11) leveraging the R library `TruncatedNormal`.

from the $N(\mathbf{x}_i^T \boldsymbol{\beta}, 1)$, and then setting $y_i = z_i \mathbb{1}(z_i > z_T)$ for every $i = 1, \dots, n$, where z_T is a pre-specified threshold to obtain the desired proportion of censored observations under the three different settings of κ . Recalling Section S1.2 in the supplementary materials, this threshold poses no difficulties in Bayesian inference since it directly enters the intercept term. To evaluate accuracy and computational efficiency at varying dimensions, these datasets are simulated for different values of $p \in \{10, 20, 50, 100, 200, 400, 800, 1200\}$. Posterior inference under the datasets produced for each combination of κ and p relies on spherical Gaussian priors $N_p(\mathbf{0}, \omega_p^2 \mathbf{I}_p)$, with $\omega_p^2 = 25 \cdot 10/p$ inducing increasing shrinkage in high dimensions. In combination with the recommended practice of standardizing the predictors to have mean 0 and standard deviation 0.5 (e.g., Chopin and Ridgway 2017), such a weakly informative prior is intended to control the overall variance of the linear predictor (see, e.g., Fasano, Durante, and Zanella 2022) so as to constrain it within a sensible range of variation for the models analyzed in this study (e.g., Gelman et al. 2008), regardless of p . This facilitates comparison across different dimensions.

Table 1 illustrates the computational benefits in sampling-based methods which can be obtained by leveraging routines that exploit directly the SUN conjugacy in Section 3. This is accomplished by comparing, for every combination of κ and p , the runtimes to obtain 5000 samples from the exact posterior distribution of $\boldsymbol{\beta}$ under both the routinely used `rstan` implementation of the No-U-Turn HMC algorithm, and the iid sampler which exploits the additive representation of the SUN posterior in (11); refer to the code at <https://github.com/niccolaoanceschi/TobitSUN> and to Chopin and Ridgway (2017) for details on the implementation of the HMC sampler in the class of models analyzed. The iid scheme leverages instead the R library `TruncatedNormal` (Botev 2017) to sample the multivariate truncated normal component in (11). As discussed in Sections 4.1 and 4.2, such a task is inherently related, in terms of implementation and computational cost, to that of evaluating the Gaussian cumulative distribution functions required to conduct posterior inference under the closed-form expressions in Section 3.2. Nonetheless, Monte Carlo inference under iid samples has the additional benefit of allowing evaluation of any functional, even beyond those derived in closed form in Section 3.2, from a single set of samples, thereby motivating our focus on sampling-based methods which allow for a more comprehensive assessment.

Consistently with related findings on probit (Durante 2019) and multinomial probit (Fasano and Durante 2022), Table 1 confirms the computational gains of iid sampling relative to

HMC in almost all the settings of κ and p , especially when p is large. In fact, while high-dimensional regimes are often challenging for HMC, under (11) p only controls the dimension of the multivariate Gaussian, which is feasible to sample from, even for a large p . As discussed in Sections 4.1–4.2, more problematic for the iid scheme is the number of censored data n_0 , which defines the dimension of the truncated normal in (11). This issue can be clearly seen in the increments of the runtimes under iid sampling when the censoring percentage grows from 15% to 85%. Nonetheless, the procedure is still competitive relative to HMC in these small-to-moderate n_0 settings. Notice also an increment in the runtime for the setting $\kappa = 0.85$ (i.e., $n_0 = 170$), when $p \approx n_0/2$. In such a regime, the method by Botev (2017) experiences low acceptance probabilities with a trend over n_0 that is reminiscent of the double-descent phenomenon in high-dimensional regression (Hastie et al. 2022). This deserves further investigations.

As clarified in Table 1, the moderate dimensions of the simulated datasets would still allow posterior inference under the closed-form solutions and iid sampling schemes presented in Sections 4.1 and 4.2. Nonetheless, as previously discussed, when n_0 grows, these procedures become computationally impractical, thus motivating also the assessment of the more scalable approximate methods presented in Section 4.3. The relevant outcomes of these performance comparisons are reported in Figures 1–2 and in Table 2, with a focus on both accuracy and scalability. In particular, Figure 1 provides insights on the accuracy of MF-VB, PFM-VB, and EP in approximating key posterior functionals of interest at varying p , and for the three different settings of κ . These quantities include the posterior mean and variance of each β_j for $j = 1, \dots, p$, along with predictive measures for the expected value of the response $\mathbb{E}[\phi(\mathbf{x}_{\text{NEW},i}^T \boldsymbol{\beta}) + (\mathbf{x}_{\text{NEW},i}^T \boldsymbol{\beta}) \Phi(\mathbf{x}_{\text{NEW},i}^T \boldsymbol{\beta}) | \mathbf{y}]$ and the probability of the censoring event $\mathbb{E}[\Phi(-\mathbf{x}_{\text{NEW},i}^T \boldsymbol{\beta}) | \mathbf{y}]$, both computed for 200 test observations whose predictors are simulated as for the original training data. For such functionals, Figure 1 displays the medians and quartiles of the absolute differences between the corresponding Monte Carlo estimates under iid sampling from the exact SUN posterior and the approximations provided by the three methods analyzed, for varying κ and p . In the first two panels, the three quartiles are computed on the p absolute differences associated with coefficients β_1, \dots, β_p , whereas in the last two panels these summaries are calculated on the 200 absolute differences for the 200 test units. For what concerns the initialization of the variational routines, we consider the default setting which sets to zero the elements of the vectors $\mathbb{E}_{q(\boldsymbol{\beta}, \bar{\mathbf{z}}_c)}(\boldsymbol{\mu}_c)$ and $\mathbb{E}_{q(\bar{\mathbf{z}}_c)}(\bar{\boldsymbol{\mu}}_c)$ in (18) and (19), respectively, for

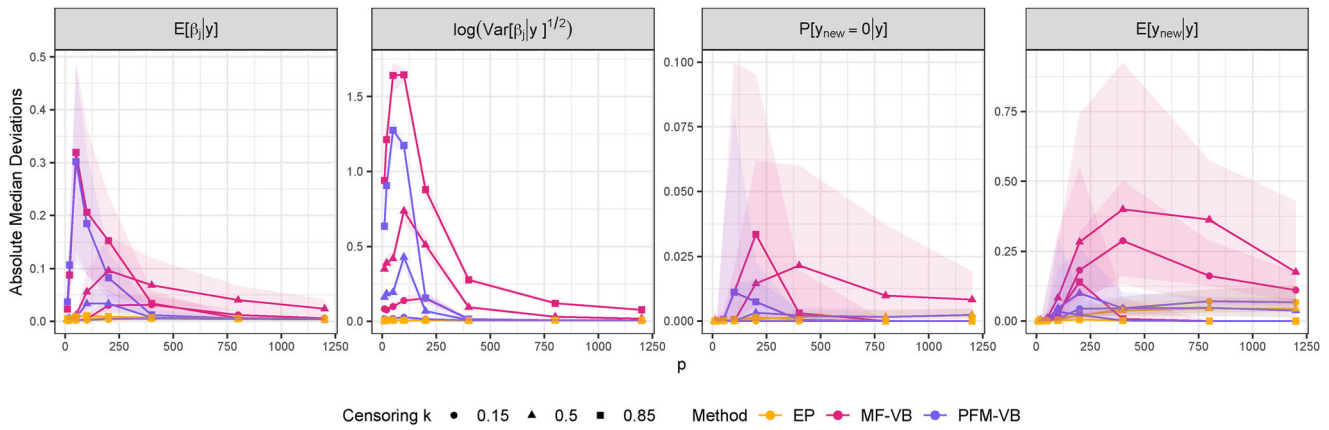


Figure 1. For four functionals of interest and different settings of $\kappa \in \{0.15, 0.50, 0.85\}$, trajectories for the median of the absolute differences, at varying p , between an accurate Monte Carlo estimate of such functionals via iid sampling from the exact SUN posterior and their approximation provided by mean-field variational Bayes (MF-VB), partially-factorized variational Bayes (PFM-VB) and expectation-propagation (EP) under tobit regression, with $n = 200$. The shaded areas correspond to the first and third quartiles computed from the absolute differences. See the online article for the color version of this figure.




Figure 2. Runtimes, in seconds, for each iteration of mean-field variational Bayes (MF-VB), partially-factorized variational Bayes (PFM-VB) and expectation-propagation (EP) in tobit regression, for $n = 200$ and different values of p and κ . The runtimes are obtained as the median over 20 repetitions. The line breaks separate the result for two alternative implementations of each routine, optimizing the computational cost in p and n_0 . For the two variational procedures, the cost per iteration is $\mathcal{O}(n_0 \cdot \min(n_0, p))$, while for EP it becomes $\mathcal{O}(n_0 p \cdot \min(n_0, p))$. See the online article for the color version of this figure.

Table 2. Total runtimes in seconds (pre-computations, total iterations and post-computations), and number of iterations (within square brackets) required to reach convergence for mean-field variational Bayes (MF-VB), partially-factorized variational Bayes (PFM-VB) and expectation-propagation (EP) under tobit regression, with $n = 200$.

Censoring	Method	p								
		10	20	50	100	200	400	800	1200	
$\kappa = 0.85$	MF-VB	0.016 [307]	0.046 [719]	0.070 [712]	0.134 [828]	0.104 [630]	0.129 [675]	0.170 [610]	0.219 [571]	
	PFM-VB	0.091 [85]	0.237 [192]	0.236 [155]	0.262 [140]	0.061 [47]	0.039 [13]	0.062 [8]	0.100 [7]	
	EP	0.014 [4]	0.018 [4]	0.068 [6]	0.270 [5]	0.175 [4]	0.252 [3]	0.327 [2]	0.488 [2]	
$\kappa = 0.50$	MF-VB	0.002 [50]	0.004 [87]	0.007 [97]	0.020 [194]	0.034 [303]	0.050 [236]	0.111 [230]	0.197 [248]	
	PFM-VB	0.009 [14]	0.016 [23]	0.017 [21]	0.024 [32]	0.021 [15]	0.033 [6]	0.086 [4]	0.165 [3]	
	EP	0.008 [4]	0.008 [3]	0.024 [4]	0.050 [4]	0.090 [5]	0.142 [4]	0.295 [4]	0.459 [4]	
$\kappa = 0.15$	MF-VB	0.001 [14]	0.001 [19]	0.001 [25]	0.003 [46]	0.019 [130]	0.043 [111]	0.126 [22]	0.253 [7]	
	PFM-VB	0.001 [5]	0.002 [6]	0.002 [7]	0.003 [8]	0.018 [7]	0.042 [3]	0.126 [3]	0.253 [3]	
	EP	0.002 [3]	0.002 [3]	0.003 [3]	0.005 [3]	0.023 [4]	0.052 [4]	0.149 [4]	0.278 [3]	

NOTE: The reported runtimes are the medians over 20 repetitions.

all $c = 1, \dots, C$. Analogously, in the case of EP we initialize the starting global approximation to the Gaussian density $l_0(\beta) = p(\beta|\bar{y}_1)$, which corresponds to setting to zero the elements of the vectors \mathbf{r}_c and matrices \mathbf{Q}_c in (23), for all $c = 1, \dots, C$. In the

absence of additional information, these initializations provide a sensible choice, which proved effective in a large number of experiments.

Consistently with Chopin and Ridgway (2017) and despite the limited theory guarantees, EP emerges as the most accurate solution in Figure 1 since its discrepancy from the Monte Carlo estimates under iid sampling is negligible in all regimes, ranging from $p < n_0$ to $p \geq n_0$. Nonetheless, as highlighted in Figure 2, the per-iteration runtimes of EP are linear in p for high dimensions, as opposed to the quadratic growth in previously-available implementations, whereas those of MF-VB and PFM-VB are essentially constant and much lower. The impact of such a scaling in the cost per iteration on the overall runtimes is relatively moderate in the regimes we focus on, as illustrated in Table 2 which displays the overall runtimes required by each routine to perform not only the total iterations, but also pre- and post-computations. Nonetheless, the higher per-iteration cost of EP becomes clearly appreciable when p exceeds a few hundreds, as seen in Figure 2 and Table 2. Therefore, from a computational perspective, the variational routines are more effective and scalable alternatives in high-dimensional settings. This is especially true for PFM-VB which, as expected from the theory in Fasano, Durante, and Zanella (2022), attains the same accuracy of EP when $p \gtrsim 2n_0$, but with a lower computational effort. On the contrary, MF-VB is not competitive with EP in terms of accuracy and does not yield remarkable improvements in runtimes relative to PFM-VB in high dimensions. Notice that, although PFM-VB and MF-VB have the same per-iteration cost, the former does not allow for a joint updating of the approximating densities for the local variables, thus requiring a for cycle in basic R implementations that yields an increment in the per-iteration runtimes under PFM-VB, due to software-related inefficiencies. This drawback can be seen from Table 2, especially in those $p < n_0$ regimes where PFM-VB requires a moderate number of iterations to reach convergence.

Table 2 also displays the number of iterations required by the three deterministic approximation procedures to reach convergence. Following common practice (e.g., Blei, Kucukelbir, and McAuliffe 2017), the convergence for the two variational methods is intended as observing a difference below a suitable threshold between the ELBO of two consecutive iterations; see Fasano, Durante, and Zanella (2022) and Consonni and Marin (2007) and the code in the GitHub repository for an expression of the ELBO under PFM-VB and MF-VB. The results presented in Table 2 correspond to a threshold of 10^{-3} on the logarithm of the ELBO. Conversely, as discussed in Section 4.3.2, EP routines do not share the same monotonicity and convergence guarantees that characterize variational routines. Therefore, the convergence of EP is commonly assessed in terms of the maximum absolute difference over the set of low dimensional parameters, characterizing the approximate Gaussian sites, between consecutive iterations (Chopin and Ridgway 2017), for which we employ the same threshold as above. Consistent with the theoretical results in Fasano, Durante, and Zanella (2022), Table 2 provides evidence on the fact that the number of iterations needed by PFM-VB to reach convergence of the ELBO decays to one as p grows to infinity, while also displaying the phenomenon reminiscent of double-descent noticed in Table 1, thus motivating further research along this line. Interestingly, the empirical

results in Table 2 also suggest that the number of iterations required by EP does not grow with p . These analyses point toward EP as a default strategy, while suggesting PFM-VB as a feasible accurate alternative in high dimensions.

6. Discussion and Future Research Directions

This review article provides a novel unified methodological and computational framework for Bayesian inference within a wide class of routinely-used regression models under a similarly broad set of prior distributions, which include the Gaussian one. Such an important gap in the literature is covered by first expressing the likelihoods associated with probit, tobit, multinomial probit and their extensions as special cases of a single formulation, and then generalizing previous findings for specific models in, for example, Durante (2019), Fasano and Durante (2022), Fasano et al. (2021), Cao, Durante, and Genton (2022), Benavoli, Azzimonti, and Piga (2020), Benavoli, Azzimonti, and Piga (2021) to prove SUN (Arellano-Valle and Azzalini 2006) conjugacy for any representation that admits such a unified likelihood. This yields general versions of past and more recent computational methods, previously proposed only for some specific members of the general class and with a focus on Gaussian priors. These include data-augmentation Gibbs samplers, iid sampling schemes, VB approximations and scalable EP implementations.

Due to the relevance of the models considered, such a review is expected to catalyze increasing interest by applied, computational and methodological researchers, and will hopefully motivate further research advancements along the directions opened by the results in Sections 2–5. For instance, the closed-form expressions in Section 3.2 for inference under the exact SUN posterior provide additional motivations to stimulate ongoing research aimed at developing accurate and fast methods to evaluate cumulative distribution functions of high-dimensional Gaussian distributions. In fact, any advancement along this direction and in sampling from multivariate truncated normals can be directly applied to conduct posterior inference via the closed-form results in Section 3.2, for increasingly larger sample sizes $\bar{n} + \bar{n}_0$, beyond small-to-moderate settings. This would be also useful for estimation of possible unknown parameters in the covariance matrices $\bar{\Sigma}_1$ and $\bar{\Sigma}_0$, via numerical maximization of the marginal likelihood $p(\mathbf{y})$ in (13). As clarified in Sections 2.1–2.4 such matrices are often parameterized by a one-dimensional or low-dimensional vector of parameters, and hence, can be effectively estimated via direct maximization of $p(\mathbf{y})$ when its evaluation is computationally practical. The availability of a closed-form expression (13) for $p(\mathbf{y})$ and of iid sampling schemes from $(\boldsymbol{\beta}|\mathbf{y})$ as in (11) can be also useful to facilitate full Bayesian inference for $\bar{\Sigma}_1$ and $\bar{\Sigma}_0$ when the associated parameters are assigned a prior. For example, leveraging $p(\mathbf{y})$ it is possible to derive collapsed Metropolis–Hastings schemes to sample from the posteriors of $\bar{\Sigma}_1$ and $\bar{\Sigma}_0$ after integrating out $\boldsymbol{\beta}$ analytically, thereby improving mixing of data-augmentation MCMC (Park and Van Dyk 2009); see also Chan and Jeliazkov (2009) for effective MCMC methods to infer $\bar{\Sigma}_0$ under identifiability constraints. These advancements are beyond the scope of this review, but provide a research

direction that is worth further exploration. Finally, it is also interesting to include hyperpriors for the scale parameters of the Gaussian or, more generally, SUN prior, that yield to scale-mixture representations inducing shrinkage in high dimensions (Carvalho, Polson, and Scott 2010). Since most of these constructions rely on conditionally Gaussian priors, the results in the present review may be useful to obtain improved theoretical and practical performance in state-of-the-art implementations of the models in Section 2 under sparse settings; see Onorati and Liseo (2022) for recent advancements along these lines, motivated by SUN conjugacy.

Although (1) already encompasses several models of interest, further generalizations of such a likelihood and of the conjugacy results in Section 3 can be considered. For instance, it is possible to extend (1) to any version of the models in Sections 2.1–2.4 that arise from censoring or rounding of the Gaussian latent utilities into a generic truncation region. As discussed in the supplementary materials, such a mechanism is directly related to the generative construction of the broader class of selection distributions (SLCT) in Arellano-Valle, Branco, and Genton (2006). Hence, following the same general reasoning, it seems natural to prove SLCT conjugacy for this broader family of likelihoods. For example, this has been done in Kowal (2022) and King and Kowal (2021) by extending the ideas in Durante (2019) and Fasano et al. (2021) to static and dynamic rounded-data situations. These generalizations can be considered to prove similar conjugacy results for any extension of (1) which incorporates truncation into a finite region. It would be also interesting to extend the recent conjugacy results under skew-elliptical link functions, such as skew- t (Branco and Dey 2001; Azzalini and Capitanio 2003; Gupta 2003), to the proposed general framework. In fact, the unified skew-elliptical distribution (Arellano-Valle and Genton 2010) has a general density expression that shares the product form of the likelihood (1), with a probability density function multiplied by a cumulative distribution function computed in some appropriate linear transformation of the parameter. This class of distributions admits as particular cases both the SUN and the unified skew- t distribution, and thus also the skew- t . Motivated by these results, Zhang et al. (in press) proved that unified skew-elliptical distributions are conjugate to probit and multinomial probit with skew-elliptical link functions, thus suggesting that this result may hold for any regression model based on skew-elliptical utilities.

As proved in a recent article by Durante, Pozza, and Szabo (2023), suitable generalizations of skew-normal distributions also provide more accurate limiting laws and skew-modal approximations for generic posterior distributions—beyond those considered in this article—with an improved convergence rate relative to the one achieved by the classical Bernstein–von Mises theorem based on limiting Gaussians.

Finally, it shall be also emphasized that the class of models discussed in Section 2 arguably encompasses one of the broadest set of formulations that appear in econometrics (Greene 2008) and social sciences (DeMaris 2004). Nonetheless, routine applications of such models under a Bayesian perspective have lagged behind the growing interest in Bayesian statistics. This review not only clarifies that the posterior distributions induced by the likelihoods of these models belong to a known class of variables, but also that such conjugacy results hold for a broader set of

priors and for various extensions of classical probit, tobit, and multinomial probit that are of direct relevance in econometrics and social sciences. This will hopefully boost the routine-use of these Bayesian models in applied research and motivate the development of even more flexible versions which still belong to likelihood (1), including, for example, random effects and graphical models (e.g., Jones et al. 2005).

Supplementary Materials

The supplementary materials contain (i) further examples of relevant regression models whose likelihood can be rewritten as in (1), (ii) proofs of Lemma 1 and Theorem 1, and (iii) a detailed discussion on the computational costs of the routines to derive the approximations in Section 4.3. Code can be found at <https://github.com/niccolaoanceschi/TobitSUN>.

Acknowledgments

We are grateful to the Editor, Associate Editor, and the two anonymous referees for the useful comments and feedbacks, which helped us in improving the preliminary version of this article.

Funding

Daniele Durante acknowledges the support from the MIUR-PRIN 2017 project 20177BRJXS.

ORCID

Augusto Fasano  <http://orcid.org/0000-0002-9734-0583>
 Daniele Durante  <http://orcid.org/0000-0002-8595-6719>
 Giacomo Zanella  <http://orcid.org/0000-0001-6727-9884>

References

- Agresti, A. (2013), *Categorical Data Analysis* (3rd ed.), Hoboken, NJ: Wiley. [1451,1452]
- Albert, J., and Chib, S. (1993), “Bayesian Analysis of Binary and Polychotomous Response Data,” *Journal of the American Statistical Association*, 88, 669–679. [1451,1452,1453,1455,1458,1459]
- (2001), “Sequential Ordinal Modeling with Applications to Survival Data,” *Biometrics*, 57, 829–836. [1451,1455,1458]
- Amemiya, T. (1984), “Tobit Models: A Survey,” *Journal of Econometrics*, 24, 3–61. [1454]
- Andrieu, C., and Doucet, A. (2002), “Particle Filtering for Partially Observed Gaussian State Space Models,” *Journal of the Royal Statistical Society, Series B*, 64, 827–836. [1451,1454]
- Arellano-Valle, R. B., and Azzalini, A. (2006), “On the Unification of Families of Skew-Normal Distributions,” *Scandinavian Journal of Statistics*, 33, 561–574. [1452,1454,1455,1456,1457,1466]
- (2021), “Some Properties of the Unified Skew-Normal Distribution,” *Statistical Papers*, 63, 461–487. [1455,1456,1457,1463]
- Arellano-Valle, R. B., Branco, M. D., and Genton, M. G. (2006), “A Unified View on Skewed Distributions Arising from Selections,” *Canadian Journal of Statistics*, 34, 581–601. [1467]
- Arellano-Valle, R. B., and Genton, M. G. (2010), “Multivariate Unified Skew-Elliptical Distributions,” *Chilean Journal of Statistics*, 1, 17–33. [1467]
- Arnold, B. C., and Beaver, R. J. (2000), “Hidden Truncation Models,” *Sankhyā: Series A*, 62, 23–35. [1455]
- Arnold, B. C., Beaver, R. J., Azzalini, A., Balakrishnan, N., Bhaumik, A., Dey, D., Cuadras, C., and Sarabia, J. M. (2002), “Skewed Multivariate Models Related to Hidden Truncation and/or Selective Reporting,” *Test*, 11, 7–54. [1455]

- Azzalini, A. (1985), "A Class of Distributions which Includes the Normal Ones," *Scandinavian Journal of Statistics*, 12, 171–178. [1455]
- Azzalini, A., and Bacchieri, A. (2010), "A Prospective Combination of Phase II and Phase III in Drug Development," *Metron*, 68, 347–369. [1455,1456,1457]
- Azzalini, A., and Capitanio, A. (1999), "Statistical Applications of the Multivariate Skew Normal Distribution," *Journal of the Royal Statistical Society, Series B*, 61, 579–602. [1455]
- (2003), "Distributions Generated by Perturbation of Symmetry with Emphasis on a Multivariate Skew t -Distribution," *Journal of the Royal Statistical Society, Series B*, 65, 367–389. [1467]
- (2013), *The Skew-Normal and Related Families*, Cambridge: Cambridge University Press. [1455,1456]
- Azzalini, A., and Dalla Valle, A. (1996), "The Multivariate Skew-Normal Distribution," *Biometrika*, 83, 715–726. [1455]
- Bazán, J. L., Bolfarine, H., and Branco, M. D. (2010), "A Framework for Skew-Probit Links in Binary Regression," *Communications in Statistics—Theory and Methods*, 39, 678–697. [1451,1452,1454]
- Benavoli, A., Azzimonti, D., and Piga, D. (2020), "Skew Gaussian Processes for Classification," *Machine Learning*, 109, 1877–1902. [1452,1454,1456,1463,1466]
- Benavoli, A., Azzimonti, D., and Piga, D. (2021), "A Unified Framework for Closed-Form Nonparametric Regression, Classification, Preference and Mixed Problems with Skew Gaussian Processes," *Machine Learning*, 110, 3095–3133. [1452,1456,1457,1463,1466]
- Bishop, C. (2006), *Pattern Recognition and Machine Learning*, New York: Springer. [1461,1462,1463]
- Blei, D. M., Kucukelbir, A., and McAuliffe, J. D. (2017), "Variational Inference: A Review for Statisticians," *Journal of the American Statistical Association*, 112, 859–877. [1452,1459,1460,1466]
- Bliss, C. I. (1934), "The Method of Probits," *Science*, 79, 38–39. [1451,1452]
- Botev, Z. I. (2017), "The Normal Law under Linear Restrictions: Simulation and Estimation via Minimax Tilting," *Journal of the Royal Statistical Society, Series B*, 88, 125–148. [1457,1458,1459,1463,1464]
- Branco, M. D., and Dey, D. K. (2001), "A General Class of Multivariate Skew-Elliptical Distributions," *Journal of Multivariate Analysis*, 79, 99–113. [1467]
- Canale, A., Pagui, E. C. K., and Scarpa, B. (2016), "Bayesian Modeling of University First-Year Students' Grades after Placement Test," *Journal of Applied Statistics*, 43, 3015–3029. [1456]
- Cao, J., Durante, D., and Genton, M. G. (2022), "Scalable Computation of Predictive Probabilities in Probit Models with Gaussian Process Priors," *Journal of Computational and Graphical Statistics*, 31, 709–720. [1452,1454,1456,1458,1461,1463,1466]
- Cao, J., Genton, M. G., Keyes, D. E., and Turkiyyah, G. M. (2019), "Hierarchical-Block Conditioning Approximations for High-Dimensional Multivariate Normal Probabilities," *Statistics and Computing*, 29, 585–598. [1457,1458]
- (2021), "Exploiting Low-Rank Covariance Structures for Computing High-Dimensional Normal and Student- t Probabilities," *Statistics and Computing*, 31, 1–16. [1457,1458]
- Carvalho, C. M., Polson, N. G., and Scott, J. G. (2010), "The Horseshoe Estimator for Sparse Signals," *Biometrika*, 97, 465–480. [1467]
- Chan, J. C.-C., and Jeliakzov, I. (2009), "MCMC Estimation of Restricted Covariance Matrices," *Journal of Computational and Graphical Statistics*, 18, 457–480. [1466]
- Chen, M.-H., Dey, D. K., and Shao, Q.-M. (1999), "A New Skewed Link Model for Dichotomous Quantal Response Data," *Journal of the American Statistical Association*, 94, 1172–1186. [1451,1452,1454]
- Chib, S. (1992), "Bayes Inference in the Tobit Censored Regression Model," *Journal of Econometrics*, 51, 79–99. [1451,1452,1453,1455,1458,1459]
- Chib, S., and Greenberg, E. (1998), "Analysis of Multivariate Probit Models," *Biometrika*, 85, 347–361. [1451,1452,1453,1455,1458]
- Chib, S., Greenberg, E., and Jeliakzov, I. (2009), "Estimation of Semiparametric Models in the Presence of Endogeneity and Sample Selection," *Journal of Computational and Graphical Statistics*, 18, 321–348. [1451]
- Chib, S., and Jeliakzov, I. (2006), "Inference in Semiparametric Dynamic Models for Binary Longitudinal Data," *Journal of the American Statistical Association*, 101, 685–700. [1454]
- Chipman, H. A., George, E. I., and McCulloch, R. E. (2010), "BART: Bayesian Additive Regression Trees," *The Annals of Applied Statistics*, 4, 266–298. [1451]
- Chopin, N., and Ridgway, J. (2017), "Leave Pima Indians Alone: Binary Regression as a Benchmark for Bayesian Computation," *Statistical Science*, 32, 64–87. [1451,1452,1455,1459,1461,1463,1464,1466]
- Consonni, G., and Marin, J.-M. (2007), "Mean-Field Variational Approximate Bayesian Inference for Latent Variable Models," *Computational Statistics & Data Analysis*, 52, 790–798. [1451,1460,1466]
- Craig, P. (2008), "A New Reconstruction of Multivariate Normal Orthant Probabilities," *Journal of the Royal Statistical Society, Series B*, 70, 227–243. [1458]
- De Oliveira, V. (2005), "Bayesian Inference and Prediction of Gaussian Random Fields based on Censored Data," *Journal of Computational and Graphical Statistics*, 14, 95–115. [1454]
- Dehaene, G., and Barthelmé, S. (2018), "Expectation Propagation in the Large Data Limit," *Journal of the Royal Statistical Society, Series B*, 80, 199–217. [1463]
- DeMaris, A. (2004), *Regression with Social Data: Modeling Continuous and Limited Response Variables*, Hoboken, NJ: Wiley. [1451,1452,1467]
- Durante, D. (2019), "Conjugate Bayes for Probit Regression via Unified Skew-Normal Distributions," *Biometrika*, 106, 765–779. [1452,1453,1456,1457,1459,1463,1464,1466,1467]
- Durante, D., Pozza, F., and Szabo, B. T. (2023), "Skewed Bernstein-von Mises Theorem and Skew-Modal Approximations," arXiv:2301.03038. [1467]
- Fasano, A., and Durante, D. (2022), "A Class of Conjugate Priors for Multinomial Probit Models which Includes the Multivariate Normal One," *Journal of Machine Learning Research*, 23, 1–26. [1452,1454,1456,1457,1459,1460,1461,1463,1464,1466]
- Fasano, A., Durante, D., and Zanella, G. (2022), "Scalable and Accurate Variational Bayes for High-Dimensional Binary Regression Models," *Biometrika*, 109, 901–919. [1452,1460,1461,1463,1464,1466]
- Fasano, A., Rebaudo, G., Durante, D., and Petrone, S. (2021), "A Closed-Form Filter for Binary Time Series," *Statistics and Computing*, 31, 1–20. [1452,1454,1456,1459,1463,1466,1467]
- Galarza Morales, C. E., Matos, L. A., Dey, D. K., and Lachos, V. H. (2022), "On Moments of Folded and Doubly Truncated Multivariate Extended Skew-Normal Distributions," *Journal of Computational and Graphical Statistics*, 31, 455–465. [1460,1461,1462]
- Gassmann, H. (2003), "Multivariate Normal Probabilities: Implementing an Old Idea of Plackett's," *Journal of Computational and Graphical Statistics*, 12, 731–752. [1458]
- Gelman, A., Jakulin, A., Pittau, M. G., and Su, Y. (2008), "A Weakly Informative Default Prior Distribution for Logistic and Other Regression Models," *The Annals of Applied Statistics*, 2, 1360–1383. [1455,1464]
- Genton, M., Keyes, D., and Turkiyyah, G. (2018), "Hierarchical Decompositions for the Computation of High-Dimensional Multivariate Normal Probabilities," *Journal of Computational and Graphical Statistics*, 27, 268–277. [1457,1458]
- Genz, A. (1992), "Numerical Computation of Multivariate Normal Probabilities," *Journal of Computational and Graphical Statistics*, 1, 141–149. [1457,1458]
- (2004), "Numerical Computation of Rectangular Bivariate and Trivariate Normal and t Probabilities," *Statistics and Computing*, 14, 251–260. [1458]
- Genz, A., and Bretz, F. (2002), "Comparison of Methods for the Computation of Multivariate t Probabilities," *Journal of Computational and Graphical Statistics*, 11, 950–971. [1458]
- (2009), *Computation of Multivariate Normal and t Probabilities*, Berlin: Springer. [1457,1458]
- Gessner, A., Kanjilal, O., and Hennig, P. (2020), "Integrals over Gaussians under Linear Domain Constraints," in *Artificial Intelligence and Statistics*, PMLR, pp. 2764–2774. [1458,1459]
- Girolami, M., and Rogers, S. (2006), "Variational Bayesian Multinomial Probit Regression with Gaussian Process Priors," *Neural Computation*, 18, 1790–1817. [1451,1460]
- González-Farías, G., Domínguez-Molina, A., and Gupta, A. K. (2004), "Additive Properties of Skew Normal Random Vectors," *Journal of Statistical Planning and Inference*, 126, 521–534. [1455]

- Greene, W. (2008), *Econometric Analysis*, New Delhi: Pearson Education India. [1451,1452,1453,1467]
- Gupta, A. (2003), "Multivariate Skew t-distribution," *Statistics: A Journal of Theoretical and Applied Statistics*, 37, 359–363. [1467]
- Gupta, A. K., González-Farías, G., and Dominguez-Molina, J. A. (2004), "A Multivariate Skew Normal Distribution," *Journal of Multivariate Analysis*, 89, 181–190. [1455]
- Hastie, T., Montanari, A., Rosset, S., and Tibshirani, R. J. (2022), "Surprises in High-Dimensional Ridgeless Least Squares Interpolation," *The Annals of Statistics*, 50, 949–986. [1464]
- Hausman, J., and Wise, D. (1978), "A Conditional Probit Model for Qualitative Choice: Discrete Decisions Recognizing Interdependence and Heterogeneous Preferences," *Econometrica*, 46, 403–426. [1451,1452,1453,1454]
- Hoffman, M. D., and Gelman, A. (2014), "The No-U-Turn Sampler: Adaptively Setting Path Lengths in Hamiltonian Monte Carlo," *Journal of Machine Learning Research*, 15, 1593–1623. [1458]
- Holmes, C. C., and Held, L. (2006), "Bayesian Auxiliary Variable Models for Binary and Multinomial Regression," *Bayesian Analysis*, 1, 145–168. [1451,1455,1458,1459,1460]
- Horrace, W. (2005), "Some Results on the Multivariate Truncated Normal Distribution," *Journal of Multivariate Analysis*, 94, 209–221. [1459,1460,1461]
- Hutton, J., and Stanghellini, E. (2011), "Modelling Bounded Health Scores with Censored Skew-Normal Distributions," *Statistics in Medicine*, 30, 368–376. [1454]
- Imai, K., and Van Dyk, D. (2005), "A Bayesian Analysis of the Multinomial Probit Model Using Marginal Data Augmentation," *Journal of Econometrics*, 124, 311–334. [1451,1455,1458]
- Johndrow, J., Dunson, D., and Lum, K. (2013), "Diagonal Orthant Multinomial Probit Models," in *Artificial Intelligence and Statistics*, PMLR, pp. 29–38. [1454]
- Johndrow, J. E., Smith, A., Pillai, N., and Dunson, D. B. (2018), "MCMC for Imbalanced Categorical Data," *Journal of the American Statistical Association*, 114, 1394–1403. [1459]
- Jones, B., Carvalho, C., Dobra, A., Hans, C., Carter, C., and West, M. (2005), "Experiments in Stochastic Computation for High-Dimensional Graphical Models," *Statistical Science*, 20, 388–400. [1467]
- King, B., and Kowal, D. R. (2021), "Warped Dynamic Linear Models for Time Series of Counts," arXiv:2110.14790. [1467]
- Kowal, D. R. (2022), "Semiparametric Discrete Data Regression with Monte Carlo Inference and Prediction," arXiv:2110.12316. [1452,1467]
- Kullback, S., and Leibler, R. (1951), "On Information and Sufficiency," *The Annals of Mathematical Statistics*, 22, 79–86. [1459]
- Kuss, M., Rasmussen, C. E., and Herbrich, R. (2005), "Assessing Approximate Inference for Binary Gaussian Process Classification," *Journal of Machine Learning Research*, 6, 1679–1704. [1451,1454,1455]
- Loaiza-Maya, R., and Nibbering, D. (2022), "Scalable Bayesian Estimation in the Multinomial Probit Model," *Journal of Business & Economic Statistics*, 40, 1678–1690. [1451]
- Loaiza-Maya, R., Smith, M. S., Nott, D. J., and Danaher, P. J. (2022), "Fast and Accurate Variational Inference for Models with Many Latent Variables," *Journal of Econometrics*, 230, 339–362. [1451]
- Manrique, A., and Shephard, N. (1998), "Simulation-based Likelihood Inference for Limited Dependent Processes," *The Econometrics Journal*, 1, 174–202. [1454]
- McCulloch, R., Polson, N., and Rossi, P. (2000), "A Bayesian Analysis of the Multinomial Probit Model with Fully Identified Parameters," *Journal of Econometrics*, 99, 173–193. [1451,1455]
- McCulloch, R., and Rossi, P. (1994), "An Exact Likelihood Analysis of the Multinomial Probit Model," *Journal of Econometrics*, 64, 207–240. [1451,1455,1458,1459]
- Minka, T. P. (2001), "Expectation Propagation for Approximate Bayesian Inference," *Proceedings of Uncertainty in Artificial Intelligence*, 17, 362–369. [1459,1461]
- Miwa, T., Hayter, A., and Kuriki, S. (2003), "The Evaluation of General Non-centred Orthant Probabilities," *Journal of the Royal Statistical Society, Series B*, 65, 223–234. [1458]
- Naveau, P., Genton, M. G., and Shen, X. (2005), "A Skewed Kalman Filter," *Journal of Multivariate Analysis*, 94, 382–400. [1454]
- Nickisch, H., and Rasmussen, C. E. (2008), "Approximations for Binary Gaussian Process Classification," *Journal of Machine Learning Research*, 9, 2035–2078. [1454]
- Nobile, A. (1998), "A Hybrid Markov Chain for the Bayesian Analysis of the Multinomial Probit Model," *Statistics and Computing*, 8, 229–242. [1451,1455]
- Onorati, P., and Liseo, B. (2022), "An Extension of the Unified Skew-Normal Family of Distributions and Application to Bayesian Binary Regression," arXiv:2209.03474. [1467]
- Pakman, A., and Paninski, L. (2014), "Exact Hamiltonian Monte Carlo for Truncated Multivariate Gaussians," *Journal of Computational and Graphical Statistics*, 23, 518–542. [1458]
- Park, T., and Van Dyk, D. A. (2009), "Partially Collapsed Gibbs Samplers: Illustrations and Applications," *Journal of Computational and Graphical Statistics*, 18, 283–305. [1466]
- Qin, Q., and Hobert, J. P. (2019), "Convergence Complexity Analysis of Albert and Chib's Algorithm for Bayesian Probit Regression," *The Annals of Statistics*, 47, 2320–2347. [1459]
- Ridgway, J. (2016), "Computation of Gaussian Orthant Probabilities in High Dimension," *Statistics and Computing*, 26, 899–916. [1458]
- Riihimäki, J., Jylänki, P., and Vehtari, A. (2014), "Nested Expectation Propagation for Gaussian Process Classification with a Multinomial Probit Likelihood," *Journal of Machine Learning Research*, 14, 75–109. [1451,1454,1455,1461]
- Roberts, G. O., and Rosenthal, J. S. (2001), "Optimal Scaling for Various Metropolis-Hastings Algorithms," *Statistical Science*, 16, 351–367. [1458]
- Rodriguez, A., and Dunson, D. (2011), "Nonparametric Bayesian Models through Probit Stick-Breaking Processes," *Bayesian Analysis*, 6, 145–178. [1451]
- Rue, H., Martino, S., and Chopin, N. (2009), "Approximate Bayesian Inference for Latent Gaussian Models by Using Integrated Nested Laplace Approximations," *Journal of the Royal Statistical Society, Series B*, 71, 319–392. [1459]
- Sahu, S. K., Dey, D. K., and Branco, M. D. (2003), "A New Class of Multivariate Skew Distributions with Applications to Bayesian Regression Models," *Canadian Journal of Statistics*, 31, 129–150. [1451,1452,1454]
- Soyer, R., and Sung, M. (2013), "Bayesian Dynamic Probit Models for the Analysis of Longitudinal Data," *Computational Statistics & Data Analysis*, 68, 388–398. [1451,1454]
- Stern, S. (1992), "A Method for Smoothing Simulated Moments of Discrete Probabilities in Multinomial Probit Models," *Econometrica*, 60, 943–952. [1451,1452,1453,1454]
- Talhouk, A., Doucet, A., and Murphy, K. (2012), "Efficient Bayesian Inference for Multivariate Probit Models with Sparse Inverse Correlation Matrices," *Journal of Computational and Graphical Statistics*, 21, 739–757. [1451]
- Tobin, J. (1958), "Estimation of Relationships for Limited Dependent Variables," *Econometrica*, 26, 24–36. [1451,1452,1454]
- Trinh, G., and Genz, A. (2015), "Bivariate Conditioning Approximations for Multivariate Normal Probabilities," *Statistics and Computing*, 25, 989–996. [1458]
- Tutz, G. (1991), "Sequential Models in Categorical Regression," *Computational Statistics & Data Analysis*, 11, 275–295. [1451,1452,1453,1454]
- Vehtari, A., Gelman, A., Sivula, T., Jylänki, P., Tran, D., Sahai, S., Blomstedt, P., Cunningham, J., Schiminovich, D., and Robert, C. (2020), "Expectation Propagation as a Way of Life: A Framework for Bayesian Inference on Partitioned Data," *Journal of Machine Learning Research*, 21, 1–53. [1461,1462,1463]
- Zellner, A. (1986), "On Assessing Prior Distributions and Bayesian Regression Analysis with g-prior Distributions," in *Bayesian Inference and Decision Techniques: Essays in Honor of Bruno de Finetti*, eds. P. Goel and A. Zellner, pp. 233–243, New York: Elsevier. [1455]
- Zhang, Z., Arellano-Valle, R. B., Genton, M. G., and Huser, R. (In press), "Tractable Bayes of Skew-Elliptical Link Models for Correlated Binary Data," *Biometrics*. [1452,1467]
- Zhang, Z., Nishimura, A., Bastide, P., Ji, X., Payne, R. P., Goulder, P., Lemey, P., and Suchard, M. A. (2021), "Large-Scale Inference of Correlation among Mixed-Type Biological Traits with Phylogenetic Multivariate Probit Models," *The Annals of Applied Statistics*, 15, 230–251. [1458]

Supplementary Materials for “Bayesian Conjugacy in Probit, Tobit, Multinomial Probit and Extensions: A Review and New Results”

Niccolò Anceschi^a, Augusto Fasano^b, Daniele Durante^{a*}, and Giacomo Zanella^a

^aDepartment of Decision Sciences and Bocconi Institute for Data Science and Analytics, Bocconi University, Milan, Italy,

^bCollegio Carlo Alberto, Turin, Italy

S1 Further extensions to skewed link functions, non-linear models and dynamic settings

Although the models discussed in Sections 2.1, 2.2 and 2.3 cover the most widely-implemented formulations in the literature, as highlighted in Sections S1.1–S1.4 several additional extensions of these representations to skewed, non-linear, dynamic and other contexts can be reframed within the likelihood in (1).

S1.1 Inclusion of skewed latent utilities

S1.1.1 Skewed extensions of linear regression and multivariate linear regression

It is possible to include skewness within the formulations in Section 2.1, while still remaining in the class of models whose likelihood can be expressed as in (1). Recalling Sahu, Dey, and Branco (2003) and Azzalini (2005), this can be done by assuming that $(y_i | \beta) \sim \text{SN}(\mathbf{x}_i^\top \beta, \sigma^2, \alpha)$, independently for each $i = 1, \dots, n$, where $\text{SN}(\mathbf{x}_i^\top \beta, \sigma^2, \alpha)$ denotes the skew-normal distribution (Azzalini, 1985) with location $\mathbf{x}_i^\top \beta$, scale σ^2 and shape parameter α . This choice implies that

$$p(\mathbf{y} | \beta) \propto \prod_{i=1}^n \phi(y_i - \mathbf{x}_i^\top \beta; \sigma^2) \Phi(\alpha(y_i - \mathbf{x}_i^\top \beta); \sigma^2) = \phi_n(\mathbf{y} - \mathbf{X}\beta; \sigma^2 \mathbf{I}_n) \Phi_n(\alpha \mathbf{y} - \alpha \mathbf{X}\beta; \sigma^2 \mathbf{I}_n), \quad (\text{S.1})$$

which coincides again with (1), when $\bar{n}_1 = \bar{n}_0 = n$, $\bar{\mathbf{y}}_1 = \mathbf{y}$, $\bar{\mathbf{y}}_0 = \alpha \mathbf{y}$, $\bar{\mathbf{X}}_1 = \mathbf{X}$, $\bar{\mathbf{X}}_0 = -\alpha \mathbf{X}$ and $\bar{\Sigma}_1 = \bar{\Sigma}_0 = \sigma^2 \mathbf{I}_n$. Inclusion of skewed responses from more elaborated distributions such as the multivariate skew-normal (Azzalini and Dalla Valle, 1996; Azzalini and Capitanio, 1999), the extended multivariate skew-normal (Arnold and Beaver, 2000; Arnold et al., 2002), the closed skew-normal family (González-Farías, Dominguez-Molina, and Gupta, 2004; Gupta, González-Farías, and Dominguez-Molina, 2004) and the SUN (Arellano-Valle and Azzalini, 2006), is also possible and yields again special cases of Equation (1); see, e.g., Canale, Pagui, and Scarpa (2016).

S1.1.2 Skewed extensions of probit, multivariate probit and multinomial probit

The inclusion of skewed link functions is possible also under probit, multivariate probit and multinomial probit models. This direction has been effectively explored by Chen, Dey, and Shao (1999) and Bazán, Bolfarine, and Branco (2010) with a main focus on basic probit models, and can be again reframed within the general formulation in (1). For example, in the context of univariate probit regression, skewness can be incorporated by replacing the Gaussian latent utilities with skew-normal ones; namely $(z_i | \beta) \sim \text{SN}(\mathbf{x}_i^\top \beta, \sigma^2, \alpha)$, independently for $i = 1, \dots, n$. As a consequence, the binary response data $y_i = \mathbb{1}(z_i > 0)$ are Bernoulli variables with probabilities $\text{pr}(y_i = 1 | \beta) = \text{pr}(z_i > 0 | \beta) \propto \Phi_2[(\mathbf{x}_i^\top \beta, 0)^\top; \text{diag}(\sigma^2, 1) + \sigma\alpha(1 + \alpha^2)^{-1/2}(\mathbf{1}_2 \mathbf{1}_2^\top - \mathbf{I}_2)]$, whose expression follows directly from the cumulative distribution function of the skew-normal; see e.g., González-Farías, Dominguez-Molina, and Gupta (2004); Arellano-Valle and Azzalini (2006); Azzalini and Bacchieri (2010); Azzalini and Capitanio (2013), and Arellano-Valle and Azzalini (2021). Leveraging the same results, it also follows that $\text{pr}(y_i = 0 | \beta) = \text{pr}(z_i < 0 | \beta) \propto \Phi_2[(-\mathbf{x}_i^\top \beta, 0)^\top; \text{diag}(\sigma^2, 1) - \sigma\alpha(1 + \alpha^2)^{-1/2}(\mathbf{1}_2 \mathbf{1}_2^\top - \mathbf{I}_2)]$. Let $\mathbf{X}_i = [(2y_i - 1)\mathbf{x}_i, 0]^\top$ and $\Sigma_i = \text{diag}(\sigma^2, 1) + (2y_i - 1)\sigma\alpha(1 + \alpha^2)^{-1/2}(\mathbf{1}_2 \mathbf{1}_2^\top - \mathbf{I}_2) = \Sigma + (2y_i - 1)\Lambda$, the joint likelihood for the binary response data can be then expressed as follows

$$p(\mathbf{y} | \beta) \propto \prod_{i=1}^n \Phi_2(\mathbf{X}_i \beta; \Sigma_i) = \Phi_{2n}(\mathbf{X}\beta; \mathbf{I}_n \otimes \Sigma + \text{diag}(2\mathbf{y} - \mathbf{1}_n) \otimes \Lambda), \quad (\text{S.2})$$

with $\mathbf{X} = (\mathbf{X}_1^\top, \dots, \mathbf{X}_n^\top)^\top$, $\Sigma = \text{diag}(\sigma^2, 1)$, and $\Lambda = \sigma\alpha(1 + \alpha^2)^{-1/2}(\mathbf{1}_2 \mathbf{1}_2^\top - \mathbf{I}_2)$. As a consequence, Equation (S.2) is again a special case of (1), after setting $\bar{n}_1 = 0$, $\bar{n}_0 = 2n$, $\bar{\mathbf{y}}_0 = \mathbf{0}$, $\bar{\mathbf{X}}_0 = \mathbf{X}$ and $\bar{\Sigma}_0 = \mathbf{I}_n \otimes \Sigma + \text{diag}(2\mathbf{y} - \mathbf{1}_n) \otimes \Lambda$. Similar derivations can be considered to incorporate skewness within multivariate and multinomial probit via multivariate skew-normal (Azzalini and Dalla Valle, 1996), closed skew-normal (González-Farías, Dominguez-Molina, and Gupta, 2004; Gupta, González-Farías, and Dominguez-Molina, 2004) or unified skew-normal

*Department of Decision Sciences and Bocconi Institute for Data Science and Analytics, Bocconi University, Italy, e-mail: daniele.durante@unibocconi.it

(Arellano-Valle and Azzalini, 2006) latent utilities. Some of these choices have not yet been explored to induce skewed link functions for multivariate and multinomial extensions of classical probit regression. Nonetheless, all these variables have cumulative distribution functions proportional to those of multivariate Gaussians, evaluated at a linear combination of β , and, hence, induce likelihoods which can be again expressed as special cases of the general framework in (1).

S1.1.3 Skewed extensions of tobit regression

As for the models presented in Sections 2.1 and 2.2, also tobit regression admits extensions to skewed contexts. This generalization has been explored, for example, by Hutton and Stanghellini (2011) who replace the Gaussian assumption for each $(z_i | \beta) \sim N(\mathbf{x}_i^\top \beta, \sigma^2)$ with the skew-normal $(z_i | \beta) \sim \text{SN}(\mathbf{x}_i^\top \beta, \sigma^2, \alpha)$, independently for every unit $i = 1, \dots, n$. Recalling the derivations for the skewed extensions of the models in Sections 2.1–2.2, this assumption implies that the contribution to the likelihood for the i -th data point is proportional to $(\phi(y_i - \mathbf{x}_i^\top \beta; \sigma^2) \Phi[\alpha(y_i - \mathbf{x}_i^\top \beta); \sigma^2])^{1(y_i > 0)} (\Phi_2[(-\mathbf{x}_i^\top \beta, 0)^\top; \text{diag}(\sigma^2, 1) - \sigma\alpha(1 + \alpha^2)^{-1/2}(\mathbf{1}_2 \mathbf{1}_2^\top - \mathbf{I}_2)])^{1(y_i = 0)}$. Therefore, letting $\text{diag}(\sigma^2, 1) = \Sigma$ and $\sigma\alpha(1 + \alpha^2)^{-1/2}(\mathbf{1}_2 \mathbf{1}_2^\top - \mathbf{I}_2) = \Lambda$, yields

$$\begin{aligned} p(\mathbf{y} | \beta) &\propto \prod_{i: y_i > 0} \phi(y_i - \mathbf{x}_i^\top \beta; \sigma^2) \Phi(\alpha(y_i - \mathbf{x}_i^\top \beta); \sigma^2) \\ &\quad \cdot \prod_{i: y_i = 0} \Phi_2[(-\mathbf{x}_i^\top \beta, 0)^\top; \Sigma - \Lambda] \\ &= \phi_{n_1}(\mathbf{y}_1 - \mathbf{X}_1 \beta; \sigma^2 \mathbf{I}_{n_1}) \\ &\quad \cdot \Phi_{n_1 + 2n_0}(\alpha[\mathbf{y}_1^\top, \mathbf{0}^\top]^\top - (\alpha \mathbf{X}_1^\top, \mathbf{X}_0^\top)^\top \beta; \Sigma_0), \end{aligned} \quad (\text{S.3})$$

where n_1 , n_0 , \mathbf{y}_1 and \mathbf{X}_1 are defined as in (7), whereas \mathbf{X}_0 is a $2n_0 \times p$ design matrix obtained by stacking $2 \times p$ row blocks $\mathbf{X}_i = (\mathbf{x}_i, \mathbf{0})^\top$ for those units with $y_i = 0$, while Σ_0 is a block-diagonal matrix with blocks $\Sigma_{0[1,1]} = \sigma^2 \mathbf{I}_{n_1}$, $\Sigma_{0[2,2]} = \mathbf{I}_{n_0} \otimes \Sigma - \mathbf{I}_{n_0} \otimes \Lambda$. Hence, to express (S.3) as a particular case of (1) it suffices to set $\bar{n}_1 = n_1$, $\bar{n}_0 = n_1 + 2n_0$, $\bar{\mathbf{y}}_1 = \mathbf{y}_1$, $\bar{\mathbf{y}}_0 = \alpha[\mathbf{y}_1^\top, \mathbf{0}^\top]^\top$, $\bar{\mathbf{X}}_1 = \mathbf{X}_1$, $\bar{\mathbf{X}}_0 = -(\alpha \mathbf{X}_1^\top, \mathbf{X}_0^\top)^\top$, $\bar{\Sigma}_1 = \sigma^2 \mathbf{I}_{n_1}$ and $\bar{\Sigma}_0 = \Sigma_0$. Recalling discussions in Sections 2.1–2.2, these derivations can be directly applied to incorporate skewness in type II–V tobit models (Amemiya, 1984), also under more general distributions which extend the original skew-normal (e.g., Arellano-Valle and Azzalini, 2006).

S1.2 Inclusion of generic thresholds

All the results presented in Sections 2.1–2.3 hold, under minor changes, when replacing the commonly-used zero threshold with a generic one z_τ , possibly varying between units. For instance, in probit regression this modification implies that $\text{pr}(y_i = 1 | \beta) = \Phi(-z_\tau + \mathbf{x}_i^\top \beta)$, thus providing the joint likelihood $\prod_{i=1}^n \Phi[(2y_i - 1)(-z_\tau + \mathbf{x}_i^\top \beta)] = \Phi_n(-z_\tau(2\mathbf{y} - \mathbf{1}_n) + \text{diag}(2\mathbf{y} - \mathbf{1}_n)\mathbf{X}\beta; \mathbf{I}_n)$,

which coincides with expression (1) after letting $\bar{n}_1 = 0$, $\bar{n}_0 = n$, $\bar{\mathbf{y}}_0 = -z_\tau(2\mathbf{y} - \mathbf{1}_n)$, $\bar{\mathbf{X}}_0 = \text{diag}(2\mathbf{y} - \mathbf{1}_n)\mathbf{X}$ and $\bar{\Sigma}_0 = \mathbf{I}_n$. Similar derivations apply to multivariate probit, multinomial probit, tobit, and their skewed extensions.

By contrast, all models relying on truncations to a finite interval of the form $[z_{1\tau}, z_{2\tau}]$ do not induce likelihoods that can be rewritten as in (1). Nonetheless, these versions are less frequent than those presented in Sections 2.1–2.3 and, as discussed in Section 6, the SUN conjugacy results presented for the general class of models whose likelihoods admit the form (1), are useful to motivate similar extensions for a generic truncation mechanism. In fact, as discussed in Arellano-Valle, Branco, and Genton (2006), the SUN family belongs itself to an even more general class of selection distributions (SLCT) whose construction rely on cumulative distribution functions evaluated at generic intervals. This result has been recently leveraged by Kowal (2022) and King and Kowal (2021) to extend the original SUN conjugacy properties presented by Durante (2019) and Fasano et al. (2021) for probit regression and its multivariate dynamic extensions, respectively, to rounded/categorical data where truncation is in finite intervals (e.g., Jeliakzov, Graves, and Kutzbach, 2008). These modifications can be extended to prove the SLCT conjugacy for generalizations of (1) which admit truncation to any finite interval.

S1.3 Inclusion of non-linear effects

Another key extension of the models presented in Sections 2.1–2.3 can be obtained by including non-linearities within the predictor. A common solution to accomplish such a goal is to replace $f(\mathbf{x}_i) = \mathbf{x}_i^\top \beta$ with the generic basis expansion $f(\mathbf{x}_i) = \mathbf{g}(\mathbf{x}_i)^\top \beta$, where $\mathbf{g}(\mathbf{x}_i) = [g_1(\mathbf{x}_i), \dots, g_k(\mathbf{x}_i)]^\top$ are pre-specified non-linear basis functions, such as splines (see e.g., Holmes and Mallick, 2001; Lang and Brezger, 2004). Including this extension within the general framework in Equation (1) poses no difficulties since it is sufficient to replicate the derivations for the models presented in Sections 2.1, 2.2 and 2.3 with $\mathbf{x}_i = (x_{i1}, \dots, x_{ip})^\top$ replaced by $\tilde{\mathbf{x}}_i = [g_1(\mathbf{x}_i), \dots, g_k(\mathbf{x}_i)]^\top$, for each $i = 1, \dots, n$.

Alternatively, one can model directly $[f(\mathbf{x}_1), \dots, f(\mathbf{x}_n)]^\top \in \mathbb{R}^n$ via a Gaussian process (e.g., Rasmussen and Williams, 2006). This direction has been commonly explored in the context of the models presented in Sections 2.1–2.3 (e.g., Kuss, Rasmussen, and Herbrich, 2005; De Oliveira, 2005; Girolami and Rogers, 2006; Nickisch and Rasmussen, 2008; Riihimäki, Jylänki, and Vehtari, 2014; Cao, Durante, and Genton, 2022; Benavoli, Azzimonti, and Piga, 2020, 2021), and can be also reframed within formulation (1). In particular, by assuming, without loss of generality, no overlap in $\mathbf{x}_1, \dots, \mathbf{x}_n$, the Gaussian process construction with mean function $m(\cdot)$ and covariance kernel $K(\cdot, \cdot)$ implies that $[f(\mathbf{x}_1), \dots, f(\mathbf{x}_n)]^\top$ is jointly distributed as a $N_n(\boldsymbol{\xi}, \boldsymbol{\Omega})$ with $\boldsymbol{\xi} = [m(\mathbf{x}_1), \dots, m(\mathbf{x}_n)]^\top$ and covariance matrix $\boldsymbol{\Omega}$ having entries $\Omega_{ii'} = K(\mathbf{x}_i, \mathbf{x}_{i'})$, for each $i = 1, \dots, n$ and $i' = 1, \dots, n$. This representation can be alternatively rewritten as $\tilde{\mathbf{X}}\beta$, where

$\beta = [f(\mathbf{x}_1), \dots, f(\mathbf{x}_n)]^\top \sim N_n(\xi, \Omega)$ and $\tilde{\mathbf{X}} = \mathbf{I}_n$. Therefore, letting $\tilde{\mathbf{x}}_i$ denote an $n \times 1$ vector with value 1 in position i and 0 elsewhere, for each $i = 1, \dots, n$, it is possible to consider Gaussian process extensions of the models in Sections 2.1–2.3, while still remaining within the general framework in (1).

S1.4 Inclusion of dynamic structure

Time-varying extensions of the models in Sections 2.1, 2.2 and 2.3 are common in the literature (e.g., Manrique and Shephard, 1998; Andrieu and Doucet, 2002; Naveau, Genton, and Shen, 2005; Chib and Jeliazkov, 2006; Soyer and Sung, 2013; Fasano et al., 2021). These extensions often appear as generalizations of the original dynamic linear model having observation equation $(\mathbf{y}_t | \beta_t) \sim N_m(\mathbf{X}_t \beta_t, \Sigma_t)$, independently for every time $t = 1, \dots, n$, and state equations $(\beta_t | \beta_{t-1}) \sim N_p(\mathbf{G}_t \beta_{t-1}, \mathbf{W}_t)$, independently for any $t = 1, \dots, n$, where \mathbf{X}_t , Σ_t , \mathbf{G}_t , and \mathbf{W}_t are known system matrices, whereas $\beta_0 \sim N_p(\mathbf{a}_0, \mathbf{P}_0)$. This building-block construction implies that the contribution to the likelihood of \mathbf{y}_t , for every $t = 1, \dots, n$, is $p(\mathbf{y}_t | \beta) = \phi_m(\mathbf{y}_t - \mathbf{X}_t \beta_t; \Sigma_t) = \phi_m(\mathbf{y}_t - \tilde{\mathbf{X}}_t \beta; \Sigma_t)$, where $\beta = (\beta_1^\top, \dots, \beta_n^\top)^\top$ and $\tilde{\mathbf{X}}_t = \mathbf{v}_t^\top \otimes \mathbf{X}_t$, with \mathbf{v}_t denoting a $n \times 1$ indicator vector having value 1 in position t and 0 elsewhere. Therefore, this representation can be directly interpreted as a particular version of the multivariate linear regression model in Equation (3) with covariance matrix possibly changing across the time units. Such a connection allows to directly recast the joint likelihood $p(\mathbf{y} | \beta)$ of $\mathbf{y} = (\mathbf{y}_1^\top, \dots, \mathbf{y}_n^\top)^\top$ within (1). Clearly, this result holds for any subsequence $\mathbf{y}_{1:t}^\top = (\mathbf{y}_1^\top, \dots, \mathbf{y}_t^\top)^\top$, with $t = 1, \dots, n$, thereby facilitating online derivation of the filtering $p(\beta_t | \mathbf{y}_{1:t})$, predictive $p(\beta_{t+1} | \mathbf{y}_{1:t})$ and smoothing $p(\beta | \mathbf{y})$ distributions via the Gaussian-Gaussian conjugacy implied by the observation and state equations (Kalman, 1960).

The above results have been recently extended by Fasano et al. (2021) for deriving the first analog of the classical Kalman filter (Kalman, 1960) within the context of multivariate dynamic probit models with Gaussian states, leveraging the SUN–probit conjugacy properties proved in Durante (2019). Recalling Fasano et al. (2021) and adapting the notation to the one in this article, the contribution to the likelihood of \mathbf{y}_t , for every time $t = 1, \dots, n$, can be expressed as $p(\mathbf{y}_t | \beta) = \Phi_m(\mathbf{B}_t \mathbf{X}_t \beta_t; \mathbf{B}_t \Sigma_t \mathbf{B}_t)$, where \mathbf{X}_t and \mathbf{B}_t are defined as in (5), with i replaced by time t , whereas Σ_t is a possibly time-varying covariance matrix among the latent utilities $(z_{t1}, \dots, z_{tm})^\top$. Recalling the derivations considered for the Gaussian dynamic setting, the expression for $p(\mathbf{y}_t | \beta)$ can be alternatively rewritten as $p(\mathbf{y}_t | \beta) = \Phi_m(\tilde{\mathbf{X}}_t \beta; \mathbf{B}_t \Sigma_t \mathbf{B}_t)$, with $\tilde{\mathbf{X}}_t = \mathbf{v}_t^\top \otimes (\mathbf{B}_t \mathbf{X}_t)$, which shows again the direct connection between this dynamic formulation and its static counterpart in (5), thereby allowing to recast the induced joint likelihood for $\mathbf{y} = (\mathbf{y}_1^\top, \dots, \mathbf{y}_n^\top)^\top$ and its subsequences $\mathbf{y}_{1:t}^\top = (\mathbf{y}_1^\top, \dots, \mathbf{y}_t^\top)^\top$, $t = 1, \dots, n$, within Equation (1).

These results clearly hold also for the dynamic extensions of

models (2) and (4), which represent the univariate versions of (3) and (5), respectively, thus simply requiring to set $m = 1$ in the above derivations. Similarly, multinomial probit (6) and tobit (7) observation equations, along with skewed extensions ((S.1), (S.2), (S.3)), can be again reframed within (1) since all these constructions are characterized by contributions to the likelihood for each time $t = 1, \dots, n$ having the same form of those associated with the statistical units $i = 1, \dots, n$ in the static counterparts of such models presented in Sections 2.1–2.3, under suitable specifications of the design and covariance matrices.

As a final remark, it is worth emphasizing that (1) naturally encompasses any combination of the models discussed in Sections 2.1–2.4. For example, if $\mathbf{y}_i = (y_{i1}, y_{i2}, y_{i3}, y_{i4})^\top$, where y_{i1} , y_{i2} , y_{i3} and y_{i4} are from models in (2), (4), (6) and (7), respectively, for each $i = 1, \dots, n$, then, leveraging the derivations in Sections 2.1–2.3, it directly follows that the joint likelihood for the vector $\mathbf{y} = (\mathbf{y}_1^\top, \dots, \mathbf{y}_n^\top)^\top$ still belongs to (1).

S2 Proofs of Lemma 1 and Theorem 1

Proof of Lemma 1. To prove Lemma 1, first notice that, by Bayes rule

$$p(\beta | \bar{\mathbf{y}}_1) \propto p(\beta) p(\bar{\mathbf{y}}_1 | \beta),$$

where $p(\bar{\mathbf{y}}_1 | \beta) = \phi_{\bar{n}_1}(\bar{\mathbf{y}}_1 - \bar{\mathbf{X}}_1 \beta; \bar{\Sigma}_1)$, whereas $p(\beta)$ is the SUN density in (8). Leveraging Gaussian-Gaussian conjugacy, it follows that the product between $\phi_{\bar{n}_1}(\bar{\mathbf{y}}_1 - \bar{\mathbf{X}}_1 \beta; \bar{\Sigma}_1)$ and the density term $\phi_{\bar{p}}(\beta - \xi; \Omega)$ in (8) is proportional to $\phi_{\bar{p}}[\beta - (\Omega^{-1} + \bar{\mathbf{X}}_1^\top \bar{\Sigma}_1^{-1} \bar{\mathbf{X}}_1)^{-1} (\Omega^{-1} \xi + \bar{\mathbf{X}}_1^\top \bar{\Sigma}_1^{-1} \bar{\mathbf{y}}_1); (\Omega^{-1} + \bar{\mathbf{X}}_1^\top \bar{\Sigma}_1^{-1} \bar{\mathbf{X}}_1)^{-1}] = \phi_{\bar{p}}(\beta - \xi_1; \Omega_1)$, where

$$\begin{aligned} \xi_1 &= (\Omega^{-1} + \bar{\mathbf{X}}_1^\top \bar{\Sigma}_1^{-1} \bar{\mathbf{X}}_1)^{-1} (\Omega^{-1} \xi + \bar{\mathbf{X}}_1^\top \bar{\Sigma}_1^{-1} \bar{\mathbf{y}}_1), \\ \Omega_1 &= (\Omega^{-1} + \bar{\mathbf{X}}_1^\top \bar{\Sigma}_1^{-1} \bar{\mathbf{X}}_1)^{-1} = \omega_1 \bar{\Omega}_1 \omega_1, \end{aligned}$$

Therefore, $p(\beta | \bar{\mathbf{y}}_1)$ is proportional to the product between this updated Gaussian density and the cumulative distribution function term $\Phi_{\bar{n}}(\gamma + \Delta^\top \bar{\Omega}^{-1} \omega^{-1} (\beta - \xi); \Gamma - \Delta^\top \bar{\Omega}^{-1} \Delta)$ of the SUN density in (8), which can be also re-expressed as $\Phi_{\bar{n}}[\mathbf{s}_1^{-1} \gamma + \mathbf{s}_1^{-1} \Delta^\top \bar{\Omega}^{-1} \omega^{-1} (\beta - \xi); \mathbf{s}_1^{-1} (\Gamma - \Delta^\top \bar{\Omega}^{-1} \Delta) \mathbf{s}_1^{-1}]$, where \mathbf{s}_1^{-1} is defined as in Lemma 1. To prove that this product yields to the SUN kernel in Lemma 1, rewrite $\mathbf{s}_1^{-1} \gamma + \mathbf{s}_1^{-1} \Delta^\top \bar{\Omega}^{-1} \omega^{-1} (\beta - \xi)$ as $\mathbf{s}_1^{-1} \gamma - \mathbf{s}_1^{-1} \Delta^\top \bar{\Omega}^{-1} \omega^{-1} \xi + \mathbf{s}_1^{-1} \Delta^\top \bar{\Omega}^{-1} \omega^{-1} \beta$, and then sum and subtract $\mathbf{s}_1^{-1} \Delta^\top \bar{\Omega}^{-1} \omega^{-1} \xi_1$ inside this expression to obtain

$$\begin{aligned} & \mathbf{s}_1^{-1} [\gamma + \Delta^\top \bar{\Omega}^{-1} \omega^{-1} (\xi_1 - \xi)] + \mathbf{s}_1^{-1} \Delta^\top \bar{\Omega}^{-1} \omega^{-1} (\beta - \xi_1) \\ &= \gamma_1 + \mathbf{s}_1^{-1} \Delta^\top \bar{\Omega}^{-1} \omega^{-1} (\beta - \xi_1) \\ &= \gamma_1 + \mathbf{s}_1^{-1} \Delta^\top \bar{\Omega}^{-1} \omega^{-1} \omega_1 \bar{\Omega}_1 \bar{\Omega}_1^{-1} \omega_1^{-1} (\beta - \xi_1) \\ &= \gamma_1 + \Delta_1^\top \bar{\Omega}_1^{-1} \omega_1^{-1} (\beta - \xi_1), \end{aligned}$$

with γ_1 and Δ_1 defined as $\gamma_1 = \mathbf{s}_1^{-1} [\gamma + \Delta^\top \bar{\Omega}^{-1} \omega^{-1} (\xi_1 - \xi)]$ and $\Delta_1 = \bar{\Omega}_1 \omega_1 \omega^{-1} \bar{\Omega}^{-1} \Delta \mathbf{s}_1^{-1}$, respectively. In order to conclude the proof, note that the correlation matrix within the cumu-

lative distribution function term can be also rewritten as

$$\begin{aligned} & \mathbf{s}_1^{-1}(\mathbf{\Gamma} - \mathbf{\Delta}^\top \bar{\mathbf{\Omega}}^{-1} \mathbf{\Delta}) \mathbf{s}_1^{-1} \\ &= \mathbf{s}_1^{-1}(\mathbf{\Gamma} - \mathbf{\Delta}^\top \bar{\mathbf{\Omega}}^{-1} \mathbf{\Delta}) \mathbf{s}_1^{-1} + \mathbf{\Delta}_1^\top \bar{\mathbf{\Omega}}_1^{-1} \mathbf{\Delta}_1 - \mathbf{\Delta}_1^\top \bar{\mathbf{\Omega}}_1^{-1} \mathbf{\Delta}_1 \\ &= \mathbf{s}_1^{-1}[\mathbf{\Gamma} + \mathbf{\Delta}^\top (\bar{\mathbf{\Omega}}^{-1} \boldsymbol{\omega}^{-1} \mathbf{\Omega}_1 \boldsymbol{\omega}^{-1} \bar{\mathbf{\Omega}}^{-1} - \bar{\mathbf{\Omega}}^{-1}) \mathbf{\Delta}] \mathbf{s}_1^{-1} \\ & \quad - \mathbf{\Delta}_1^\top \bar{\mathbf{\Omega}}_1^{-1} \mathbf{\Delta}_1, \end{aligned}$$

which corresponds to $\mathbf{\Gamma}_1 - \mathbf{\Delta}_1^\top \bar{\mathbf{\Omega}}_1^{-1} \mathbf{\Delta}_1$, with $\mathbf{\Gamma}_1$ defined as in Lemma 1. This proves that the kernel $p(\boldsymbol{\beta})p(\bar{\mathbf{y}}_1 | \boldsymbol{\beta})$ of the posterior coincides with that of a SUN having parameters $\boldsymbol{\xi}_1, \mathbf{\Omega}_1, \mathbf{\Delta}_1, \boldsymbol{\gamma}_1$ and $\mathbf{\Gamma}_1$ specified as in Lemma 1. \square

Proof of Theorem 1. The proof of Theorem 1 simply requires to combine Lemma 1 with an adaptation of Corollary 4 in Durante (2019). In particular, by direct application of the Bayes rule, it follows that

$$p(\boldsymbol{\beta} | \mathbf{y}) \propto p(\boldsymbol{\beta})p(\mathbf{y} | \boldsymbol{\beta}) \propto [p(\boldsymbol{\beta})p(\bar{\mathbf{y}}_1 | \boldsymbol{\beta})]p(\bar{\mathbf{y}}_0 | \boldsymbol{\beta}).$$

Therefore, the posterior distribution $p(\boldsymbol{\beta} | \mathbf{y})$ can be obtained by first updating the SUN prior $p(\boldsymbol{\beta})$ with $p(\bar{\mathbf{y}}_1 | \boldsymbol{\beta}) \propto \phi_{\bar{n}_1}(\bar{\mathbf{y}}_1 - \bar{\mathbf{X}}_1 \boldsymbol{\beta}; \bar{\mathbf{\Sigma}}_1)$, and then use such a conditional density $p(\boldsymbol{\beta} | \bar{\mathbf{y}}_1)$ as an intermediate prior to be updated with the likelihood $p(\bar{\mathbf{y}}_0 | \boldsymbol{\beta}) \propto \Phi_{\bar{n}_0}(\bar{\mathbf{y}}_0 + \bar{\mathbf{X}}_0 \boldsymbol{\beta}; \bar{\mathbf{\Sigma}}_0)$ of $\bar{\mathbf{y}}_0$ for obtaining the final posterior. By direct application of Lemma 1, it follows that $p(\boldsymbol{\beta} | \bar{\mathbf{y}}_1)$ is the density of the $\text{SUN}_{\bar{p}, \bar{n}}(\boldsymbol{\xi}_1, \mathbf{\Omega}_1, \mathbf{\Delta}_1, \boldsymbol{\gamma}_1, \mathbf{\Gamma}_1)$ with parameters defined as in Lemma 1. Therefore, to conclude the proof, it is sufficient to prove that the updating of this intermediate prior with the likelihood $p(\bar{\mathbf{y}}_0 | \boldsymbol{\beta}) \propto \Phi_{\bar{n}_0}(\bar{\mathbf{y}}_0 + \bar{\mathbf{X}}_0 \boldsymbol{\beta}; \bar{\mathbf{\Sigma}}_0)$ for $\bar{\mathbf{y}}_0$ yields again to a $\text{SUN}_{\bar{p}, \bar{n} + \bar{n}_0}(\boldsymbol{\xi}_{\text{post}}, \mathbf{\Omega}_{\text{post}}, \mathbf{\Delta}_{\text{post}}, \boldsymbol{\gamma}_{\text{post}}, \mathbf{\Gamma}_{\text{post}})$ with $\boldsymbol{\xi}_{\text{post}}, \mathbf{\Omega}_{\text{post}}, \mathbf{\Delta}_{\text{post}}, \boldsymbol{\gamma}_{\text{post}}$ and $\mathbf{\Gamma}_{\text{post}}$ defined as in Theorem 1. This result follows directly from an adaptation of Corollary 4 in Durante (2019); refer also to Theorem 1 in Fasano and Durante (2022). In particular, replacing \mathbf{D} with $\bar{\mathbf{X}}_0$ and \mathbf{I}_n with $\bar{\mathbf{\Sigma}}_0$ in Corollary 4 by Durante (2019), under a $\text{SUN}_{\bar{p}, \bar{n}}(\boldsymbol{\xi}_1, \mathbf{\Omega}_1, \mathbf{\Delta}_1, \boldsymbol{\gamma}_1, \mathbf{\Gamma}_1)$ prior, yields to the expressions for $\boldsymbol{\xi}_{\text{post}}, \mathbf{\Omega}_{\text{post}}, \mathbf{\Delta}_{\text{post}}$ and $\mathbf{\Gamma}_{\text{post}}$ in Theorem 1. Inclusion of the offset $\bar{\mathbf{y}}_0$ in the proof of Corollary 4 by Durante (2019) poses no difficulties since it directly enters the SUN truncation parameter, thereby providing the expression for $\boldsymbol{\gamma}_{\text{post}}$ in Theorem 1. \square

S3 Computational costs

To further extend the analysis of the approximate schemes in Sections 4.3.1–4.3.2, we discuss the associated cost per-iteration focusing, for the ease of notation, on the classical probit regression $\prod_{i=1}^n \Phi(\mathbf{x}_i^\top \boldsymbol{\beta})^{\mathbb{1}(y_i=1)} [1 - \Phi(\mathbf{x}_i^\top \boldsymbol{\beta})]^{\mathbb{1}(y_i=0)}$ as in (4), with a spherical Gaussian prior $p(\boldsymbol{\beta}) = \phi_p(\boldsymbol{\beta}; \omega^2 \mathbf{I}_p)$. Note that, as discussed in Section 2.2, $n = \bar{n}_0$ and $p = \bar{p}$ when such a model is written as a special case of likelihood (1). Besides providing one of the most widely implemented formulations within the class of models whose likelihood can be expressed as in (1), this choice is also motivated by the fact that detailed costs per-iteration of effective

MF-VB and PFM-VB implementations have been already derived in Fasano, Durante, and Zanella (2022) under probit regression with a Gaussian prior. Moreover, it gives the opportunity to show that currently-reported per-iteration costs of EP for the same class of models and priors (Chopin and Ridgway, 2017) can be further reduced, thus making also EP more scalable to high dimensions.

For deriving the costs of MF-VB, PFM-VB and EP it shall be emphasized that, in probit regression, such approximations rely on $c = 1, \dots, n$ and, therefore, $n_c = \bar{n}_c = 1$ for any c . Under EP (Chopin and Ridgway, 2017) this choice is a common practice which follows from the factorized form of the likelihood, whereas in MF-VB (Consonni and Marin, 2007) it is implied by the formulation of the optimization problem and is a direct consequence of the conditional independence among the unit-specific latent utilities. Instead, for PFM-VB (Fasano, Durante, and Zanella, 2022) such a setting is not enforced. Nonetheless, it provides a convenient specification which is in line with MF-VB and EP solutions, and also facilitates posterior inference by only requiring to deal with univariate truncated normals.

Under the aforementioned settings, Appendix A of Fasano, Durante, and Zanella (2022) provides a detailed discussion of the per-iteration cost for both MF-VB and also PFM-VB, which is $\mathcal{O}(n \cdot \min\{n, p\}) = \mathcal{O}(\bar{n}_0 \cdot \min\{\bar{n}_0, \bar{p}\})$, after suitable matrix precomputations before running the CAVI routine. Since the mean and variance of univariate truncated normals can be accurately computed at $\mathcal{O}(1)$ cost under standard algorithms (e.g., Botev, 2017), the most intensive computations in the CAVI routines for MF-VB and PFM-VB correspond to the matrix multiplication operations. These steps can be efficiently implemented by exploiting recursive formulas when updating each univariate truncated normal approximating density conditioned the most recent estimate of the others, in PFM-VB, or of the $\boldsymbol{\beta}$ parameters, in MF-VB, thereby leading to an overall per-iteration cost that is either linear or sublinear in $p = \bar{p}$.

As for EP, the currently reported per-iteration cost in probit regression with spherical Gaussian priors is $\mathcal{O}(np^2) = \mathcal{O}(\bar{n}_0 \bar{p}^2)$ (Chopin and Ridgway, 2017), after suitable precomputations as in MF-VB and PFM-VB. Intuitively, this increased complexity is due to the fact that, unlike for MF-VB and PFM-VB, not only the expectations but also the $p \times p$ covariance matrices must be updated and then inverted at every site c , for each $c = 1, \dots, n$. Although the specific form of such matrices allows to reduce the common cubic cost into a quadratic one via the application of the Woodbury’s formula to avoid direct matrix inversion (Chopin and Ridgway, 2017), an $\mathcal{O}(np^2) = \mathcal{O}(\bar{n}_0 \bar{p}^2)$ cost can be still computationally impractical in the large $p = \bar{p}$ setting. In fact, as mentioned in the final discussion by Chopin and Ridgway (2017), even the state-of-the-art implementations of EP are often computationally challenging when p exceeds one thousand. This point is also confirmed in the empirical studies of Fasano, Durante, and Zanella (2022), where the EP implementation within the R package EPGLM by Chopin and Ridgway (2017) requires more than six

hours to reach convergence in a high-dimensional Alzheimer’s application with $p = \bar{p} = 9036$ and $n = \bar{n}_0 = 300$. Notably, as a further contribution of the present article, it shall be emphasized that a more scalable EP implementation with per-iteration cost $\mathcal{O}(np \cdot \min\{n, p\}) = \mathcal{O}(\bar{n}_0 \bar{p} \cdot \min\{\bar{n}_0, \bar{p}\})$ can be actually derived by leveraging similar results considered in [Fasano, Durante, and Zanella \(2022\)](#) for obtaining efficient implementations of MF-VB and PFM-VB. In particular, this new EP implementation exploits the fact that, under the same reformulation via Woodbury’s identity of [Chopin and Ridgway \(2017\)](#), the site updates do not necessarily require direct computation of the aforementioned $p \times p$ matrices, since such quantities enter via the inner product with the $p \times n$ design matrices $\bar{\mathbf{X}}_0^\top$. Hence, when $p = \bar{p}$ is large, it is more convenient to update this product directly, without storing, updating or multiplying any $p \times p$ matrix. This yields to an $\mathcal{O}(np) = \mathcal{O}(\bar{n}_0 \bar{p})$ cost for each site, and to an overall cost for the $n = \bar{n}_0$ site updates of $\mathcal{O}(n^2 p) = \mathcal{O}(\bar{n}_0^2 \bar{p})$. In high dimensional settings, when $p \gg n$, this linear cost in $p = \bar{p}$ yields to massive computational gains relative to the original $\mathcal{O}(np^2) = \mathcal{O}(\bar{n}_0 \bar{p}^2)$ cost of EPGLM in [Chopin and Ridgway \(2017\)](#). For example, by considering the proposed more scalable implementation within the high-dimensional Alzheimer’s application yields to an overall runtime of less than one minute, which is orders of magnitude lower than EPGLM ([Chopin and Ridgway, 2017](#)) that requires, instead, more than six hours. When, instead, $n \gg p$, the linear cost in $n = \bar{n}_0$ of the EPGLM ensures effective implementations. Combining these two scenarios yields to an overall per-iteration cost of $\mathcal{O}(np \cdot \min\{n, p\}) = \mathcal{O}(\bar{n}_0 \bar{p} \cdot \min\{\bar{n}_0, \bar{p}\})$, which is linear in the higher between $n = \bar{n}_0$ and $p = \bar{p}$. To the best of our knowledge, this is the first implementation of EP available in the literature to achieve such a computational efficiency. Despite this, standard EP remains more computationally demanding than MF-VB and PFM-VB since, as discussed above, also the $p \times p$ matrices need to be updated at each step of the EP routine, either directly or implicitly via the product with $\bar{\mathbf{X}}_0^\top$.

The above reasoning can be directly applied in order to highlight a similar dependence on sample size and number of predictors in the per-iteration cost of effective MF-VB, PFM-VB and EP implementations under the whole class of models and priors in Sections 2 and 3 — as long as n_c and \bar{n}_c are sufficiently small to allow the calculations of the moments for the associated multivariate truncated normals at a negligible cost compared to the one of the matrix operations. This result is illustrated in empirical studies in Section 5, with a focus on tobit regression.

References

- Amemiya, T. (1984), “Tobit models: A survey,” *Journal of Econometrics*, 24, 3–61.
- Andrieu, C., and Doucet, A. (2002), “Particle filtering for partially observed Gaussian state space models,” *Journal of the Royal Statistical Society: Series B*, 64, 827–836.
- Arellano-Valle, R. B., and Azzalini, A. (2006), “On the unification of families of skew-normal distributions,” *Scandinavian Journal of Statistics*, 33, 561–574.
- Arellano-Valle, R. B.— (2021), “Some properties of the unified skew-normal distribution,” *Statistical Papers*, 63, 461–487.
- Arellano-Valle, R. B., Branco, M. D., and Genton, M. G. (2006), “A unified view on skewed distributions arising from selections,” *Canadian Journal of Statistics*, 34, 581–601.
- Arnold, B. C., and Beaver, R. J. (2000), “Hidden truncation models,” *Sankhyā: Series A*, 62, 23–35.
- Arnold, B. C., Beaver, R. J., Azzalini, A., Balakrishnan, N., Bhaumik, A., Dey, D., Cuadras, C., and Sarabia, J. M. (2002), “Skewed multivariate models related to hidden truncation and/or selective reporting,” *Test*, 11, 7–54.
- Azzalini, A. (1985), “A class of distributions which includes the normal ones,” *Scandinavian Journal of Statistics*, 12, 171–178.
- (2005), “The skew-normal distribution and related multivariate families,” *Scandinavian Journal of Statistics*, 32, 159–188.
- Azzalini, A., and Bacchieri, A. (2010), “A prospective combination of phase II and phase III in drug development,” *Metron*, 68, 347–369.
- Azzalini, A., and Capitanio, A. (1999), “Statistical applications of the multivariate skew normal distribution,” *Journal of the Royal Statistical Society: Series B*, 61, 579–602.
- Azzalini, A.— (2013), *The Skew-Normal and Related Families*, Cambridge University Press.
- Azzalini, A., and Dalla Valle, A. (1996), “The multivariate skew-normal distribution,” *Biometrika*, 83, 715–726.
- Bazán, J. L., Bolfarine, H., and Branco, M. D. (2010), “A framework for skew-probit links in binary regression,” *Communications in Statistics—Theory and Methods*, 39, 678–697.
- Benavoli, A., Azzimonti, D., and Piga, D. (2020), “Skew Gaussian processes for classification,” *Machine Learning*, 109, 1877–1902.
- Benavoli, A., Azzimonti, D., and Piga, D. (2021), “A unified framework for closed-form nonparametric regression, classification, preference and mixed problems with skew Gaussian processes,” *Machine Learning*, 110, 3095–3133.
- Botev, Z. I. (2017), “The normal law under linear restrictions: Simulation and estimation via minimax tilting,” *Journal of the Royal Statistical Society: Series B*, 88, 125–148.
- Canale, A., Pagui, E. C. K., and Scarpa, B. (2016), “Bayesian modeling of university first-year students’ grades after placement test,” *Journal of Applied Statistics*, 43, 3015–3029.
- Cao, J., Durante, D., and Genton, M. G. (2022), “Scalable computation of predictive probabilities in probit models with Gaussian process priors,” *Journal of Computational and Graphical Statistics*, 31, 709–720.

- Chen, M.-H., Dey, D. K., and Shao, Q.-M. (1999), “A new skewed link model for dichotomous quantal response data,” *Journal of the American Statistical Association*, 94, 1172–1186.
- Chib, S., and Jeliazkov, I. (2006), “Inference in semiparametric dynamic models for binary longitudinal data,” *Journal of the American Statistical Association*, 101, 685–700.
- Chopin, N., and Ridgway, J. (2017), “Leave Pima Indians alone: Binary regression as a benchmark for Bayesian computation,” *Statistical Science*, 32, 64–87.
- Consonni, G., and Marin, J.-M. (2007), “Mean-field variational approximate Bayesian inference for latent variable models,” *Computational Statistics & Data Analysis*, 52, 790–798.
- De Oliveira, V. (2005), “Bayesian inference and prediction of Gaussian random fields based on censored data,” *Journal of Computational and Graphical Statistics*, 14, 95–115.
- Durante, D. (2019), “Conjugate Bayes for probit regression via unified skew-normal distributions,” *Biometrika*, 106, 765–779.
- Fasano, A., and Durante, D. (2022), “A class of conjugate priors for multinomial probit models which includes the multivariate normal one,” *Journal of Machine Learning Research*, 23, 1–26.
- Fasano, A., Durante, D., and Zanella, G. (2022), “Scalable and accurate variational Bayes for high-dimensional binary regression models,” *Biometrika*, 109, 901–919.
- Fasano, A., Rebaudo, G., Durante, D., and Petrone, S. (2021), “A closed-form filter for binary time series,” *Statistics and Computing*, 31, 1–20.
- Girolami, M., and Rogers, S. (2006), “Variational Bayesian multinomial probit regression with Gaussian process priors,” *Neural Computation*, 18, 1790–1817.
- González-Farías, G., Dominguez-Molina, A., and Gupta, A. K. (2004), “Additive properties of skew normal random vectors,” *Journal of Statistical Planning and Inference*, 126, 521–534.
- Gupta, A. K., González-Farías, G., and Dominguez-Molina, J. A. (2004), “A multivariate skew normal distribution,” *Journal of Multivariate Analysis*, 89, 181–190.
- Holmes, C., and Mallick, B. (2001), “Bayesian regression with multivariate linear splines,” *Journal of the Royal Statistical Society: Series B*, 63, 3–17.
- Hutton, J., and Stanghellini, E. (2011), “Modelling bounded health scores with censored skew-normal distributions,” *Statistics in Medicine*, 30, 368–376.
- Jeliazkov, I., Graves, J., and Kutzbach, M. (2008), “Fitting and comparison of models for multivariate ordinal outcomes,” in *Bayesian Econometrics*, vol. 23, pp. 115–156.
- Kalman, R. E. (1960), “A new approach to linear filtering and prediction problems,” *Journal of Basic Engineering*, 82, 35–45.
- King, B., and Kowal, D. R. (2021), “Warped dynamic linear models for time series of counts,” *arXiv:2110.14790*.
- Kowal, D. R. (2022), “Semiparametric discrete data regression with Monte Carlo inference and prediction,” *arXiv:2110.12316*.
- Kuss, M., Rasmussen, C. E., and Herbrich, R. (2005), “Assessing approximate inference for binary Gaussian process classification,” *Journal of Machine Learning Research*, 6, 1679–1704.
- Lang, S., and Brezger, A. (2004), “Bayesian P-splines,” *Journal of Computational and Graphical Statistics*, 13, 183–212.
- Manrique, A., and Shephard, N. (1998), “Simulation-based likelihood inference for limited dependent processes,” *The Econometrics Journal*, 1, 174–202.
- Naveau, P., Genton, M. G., and Shen, X. (2005), “A skewed Kalman filter,” *Journal of Multivariate Analysis*, 94, 382–400.
- Nickisch, H., and Rasmussen, C. E. (2008), “Approximations for binary Gaussian process classification,” *Journal of Machine Learning Research*, 9, 2035–2078.
- Rasmussen, C. E., and Williams, C. K. I. (2006), *Gaussian Processes for Machine Learning*, MIT Press.
- Riihimäki, J., Jylänki, P., and Vehtari, A. (2014), “Nested expectation propagation for Gaussian process classification with a multinomial probit likelihood,” *Journal of Machine Learning Research*, 14, 75–109.
- Sahu, S. K., Dey, D. K., and Branco, M. D. (2003), “A new class of multivariate skew distributions with applications to Bayesian regression models,” *Canadian Journal of Statistics*, 31, 129–150.
- Soyer, R., and Sung, M. (2013), “Bayesian dynamic probit models for the analysis of longitudinal data,” *Computational Statistics & Data Analysis*, 68, 388–398.